Oracle® Private Cloud Appliance

Administrator's Guide for Release 2.0

E56837-04
November 2015

Oracle Legal Notices

Copyright © 2013, 2015, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Table of Contents

Preface vii
1 Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance 1
1.1 What is Oracle Private Cloud Appliance 1
1.2 Hardware Components 2
1.2.1 Management Nodes 3
1.2.2 Compute Nodes 4
1.2.3 Storage Appliance 5
1.2.4 Network Infrastructure 6
1.3 Software Components 8
1.3.1 Oracle Private Cloud Appliance Dashboard 8
1.3.2 Password Manager (Wallet) 8
1.3.3 Oracle VM Manager 9
1.3.4 Operating Systems 9
1.3.5 Databases 9
1.3.6 Oracle Private Cloud Appliance Management Software 12
1.3.7 Oracle Private Cloud Appliance Diagnostics Tool 13
1.4 Provisioning and Orchestration 14
1.4.1 Appliance Management Initialization 14
1.4.2 Compute Node Discovery and Provisioning 14
1.4.3 Server Pool Readiness 15
1.5 High Availability 16
1.6 Oracle Private Cloud Appliance Backup 17
1.7 Oracle Private Cloud Appliance Software Update 18
2 Monitoring and Managing Oracle Private Cloud Appliance 21
2.1 Connecting and Logging in to the Oracle Private Cloud Appliance Dashboard 22
2.2 Hardware View 22
2.3 Network View 25
2.4 Network Settings 25
2.5 Authentication 27
2.6 Update 28
2.6.1 Mandatory Updates 29
2.6.2 Updating the Oracle PCA Controller Software 30
2.6.3 Upgrading the Virtualization Platform 33
2.6.4 Upgrading Component Firmware 35
3 The Oracle Private Cloud Appliance Command Line Interface (CLI) 37
3.1 CLI Usage 37
3.1.1 Interactive Mode 38
3.1.2 Single-command Mode 39
3.1.3 Controlling CLI Output 40
3.1.4 Internal CLI Help 42
3.2 CLI Commands 43
3.2.1 backup 43
3.2.2 diagnose 44
3.2.3 delete 45
3.2.4 list 46
3.2.5 reprovision 50
3.2.6 rerun 50
3.2.7 set 51
3.2.8 show 52
3.2.9 start 56
3.2.10 stop 56

4 Managing the Oracle VM Virtual Infrastructure 57
4.1 Guidelines and Limitations 58
4.2 Logging in to the Oracle VM Manager Web UI 58
4.3 Monitoring Health and Performance in Oracle VM 59
4.4 Creating and Managing Virtual Machines 59
4.5 Managing Virtual Machine Resources 62
4.6 Configuring Network Resources for Virtual Machines 64
4.7 Viewing and Managing Storage Resources 66
4.8 Tagging of Resources in Oracle VM Manager 67
4.9 Managing Jobs and Events 67
5 Automating Oracle Private Cloud Appliance Service Requests 69
5.1 Understanding Oracle Auto Service Request (ASR) 69
5.2 ASR Prerequisites 70
5.3 Installing ASR Manager 71
5.4 Activating ASR Assets 75
5.4.1 Activating Compute Nodes 75
5.4.2 Activating InfiniBand Switches 76
5.4.3 Activating ASR on the ZFS Storage Appliance 78
6 Servicing Oracle Private Cloud Appliance Components 81
6.1 Replaceable Components 81
6.1.1 Rack Components 82
6.1.2 Oracle Server X5-2 Components 83
6.1.3 Sun Server X4-2 Components 84
6.1.4 Sun Server X3-2 Components 84
6.1.5 Oracle ZFS Storage Appliance ZS3-ES Components 85
6.1.6 Sun ZFS Storage Appliance 7320 Components 86
6.1.7 Oracle Switch ES1-24 Components 88
6.1.8 NM2-36P Sun Datacenter InfiniBand Expansion Switch Components 88
6.1.9 Oracle Fabric Interconnect F1-15 Components 89
6.2 Preparing Oracle Private Cloud Appliance for Service 89
6.3 Servicing the Oracle Private Cloud Appliance Rack System 90
6.3.1 Powering Down Oracle Private Cloud Appliance (When Required) 90
6.3.2 Service Procedures for Rack System Components 91
6.4 Servicing an Oracle Server X5-2 92
6.4.1 Powering Down Oracle Server X5-2 for Service (When Required) 92
6.4.2 Service Procedures for Oracle Server X5-2 Components 93
6.5 Servicing a Sun Server X4-2 94
6.5.1 Powering Down Sun Server X4-2 for Service (When Required) 94
6.5.2 Service Procedures for Sun Server X4-2 Components 96
6.6 Servicing a Sun Server X3-2 97
6.6.1 Powering Down Sun Server X3-2 for Service (When Required) 97
6.6.2 Service Procedures for Sun Server X3-2 Components 98
6.7 Servicing the Oracle ZFS Storage Appliance ZS3-ES 99
6.7.1 Powering Down the Oracle ZFS Storage Appliance ZS3-ES for Service (When Required) 99
6.7.2 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components 101
6.8 Servicing the Sun ZFS Storage Appliance 7320 102
6.8.1 Powering Down the Sun ZFS Storage Appliance 7320 for Service (When Required) ... 103
6.8.2 Service Procedures for Sun ZFS Storage Appliance 7320 Components 104
6.9 Servicing an Oracle Switch ES1-24 105
6.9.1 Powering Down the Oracle Switch ES1-24 for Service (When Required) 105
6.9.2 Service Procedures for Oracle Switch ES1-24 Components 105
6.10 Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch 106

6.10.1 Powering Down the NM2-36P Sun Datacenter InfiniBand Expansion Switch for Service (When Required) 106
6.10.2 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components 106
6.11 Servicing an Oracle Fabric Interconnect F1-15 107
6.11.1 Powering Down the Oracle Fabric Interconnect F1-15 for Service (When Required) 107
6.11.2 Service Procedures for Oracle Fabric Interconnect F1-15 Components 107
7 Troubleshooting 111
7.1 Adding Proxy Settings for Oracle Private Cloud Appliance Updates 111
7.2 Setting the Oracle Private Cloud Appliance Logging Parameters 112
7.3 Configuring Data Center Switches for VLAN Traffic 112
7.4 Enabling Support Tools on Systems Provisioned with Software Release 1.0.2 113
7.5 Verifying and Re-applying Oracle VM Manager Tuning after Software Update 114
7.6 Upgrading Existing Compute Node Configuration from Release 1.0.2 115
7.7 Enabling Compute Node IPMI in Oracle VM Manager 118
7.8 Enabling SNMP Server Monitoring 118
7.9 Changing Oracle WebLogic Server Passwords 120
7.10 Replacing Default Passwords Manually 121
7.11 Using a Custom CA Certificate for SSL Encryption 123
7.12 A Compute Node Fails to Complete Provisioning 123
7.13 Oracle VM Server Pool Is Offline After Network Services Restart 125
Index 127

Preface

This document is part of the documentation library for Oracle Private Cloud Appliance (PCA) Release 2.0, which is available at:

http://docs.oracle.com/cd/E55813_01.

The documentation library consists of the following items:

Oracle Private Cloud Appliance Release Notes
The release notes provide a summary of the new features, changes, fixed bugs and known issues in Oracle PCA.

Oracle Private Cloud Appliance Licensing Information User Manual
The licensing information user manual provides information about the various product licenses applicable to the use of Oracle PCA.

Oracle Private Cloud Appliance Installation Guide
The installation guide provides detailed instructions to prepare the installation site and install Oracle PCA. It also includes the procedures to install additional compute nodes, and to connect and configure external storage components.

Oracle Private Cloud Appliance Safety and Compliance Guide
The safety and compliance guide is a supplemental guide to the safety aspects of Oracle PCA. It conforms to Compliance Model No. OCA1A.

Oracle Private Cloud Appliance Administrator's Guide
The administrator's guide provides instructions for using the management software. It is a comprehensive guide to how to configure, monitor and administer Oracle PCA.

Oracle Private Cloud Appliance Quick Start Poster
The quick start poster provides a step-by-step description of the hardware installation and initial software configuration of Oracle PCA. A printed quick start poster is shipped with each Oracle PCA base rack, and is intended for data center operators and administrators who are new to the product.

The quick start poster is also available in the documentation library as an HTML guide, which contains alternate text for ADA 508 compliance.

Oracle Private Cloud Appliance Expansion Node Setup Poster
The expansion node setup poster provides a step-by-step description of the installation procedure for an Oracle PCA expansion node. A printed expansion node setup poster is shipped with each Oracle PCA expansion node.

The expansion node setup poster is also available in the documentation library as an HTML guide, which contains alternate text for ADA 508 compliance.

Audience

The Oracle Private Cloud Appliance (PCA) documentation is written for technicians, authorized service providers, data center operators and system administrators who want to install, configure and maintain a private cloud environment in order to deploy virtual machines for users. It is assumed that readers have experience installing and troubleshooting hardware, are familiar with web and virtualization technologies and have a general understanding of operating systems such as UNIX (including Linux) and Windows.

The Oracle PCA makes use of Oracle Linux and Oracle Solaris operating systems within its component configuration. It is advisable that administrators have experience of these operating systems at the very

least. Oracle PCA is capable of running virtual machines with a variety of operating systems including Oracle Solaris and other UNIXes, Linux and Microsoft Windows. The selection of operating systems deployed in guests on Oracle PCA determines the requirements of your administrative knowledge.

Related Documentation

Additional Oracle components may be included with Oracle PCA depending on configuration. The documentation for such additional components is available as follows:

- Sun Rack II 1042 and 1242

  http://docs.oracle.com/cd/E19844-01/index.html

- Oracle Server X5-2

  http://docs.oracle.com/cd/E41059_01/index.html

- Sun Server X4-2

  http://docs.oracle.com/cd/E36975_01/index.html

- Sun Server X3-2

  http://docs.oracle.com/cd/E22368_01/index.html

- Oracle ZFS Storage Appliance ZS3-ES

  http://docs.oracle.com/cd/E37831_01/index.html

- Sun ZFS Storage Appliance 7320

  http://docs.oracle.com/cd/E28317_01/index.html

- Oracle Switch ES1-24

  http://docs.oracle.com/cd/E39109_01/index.html

- NM2-36P Sun Datacenter InfiniBand Expansion Switch

  http://docs.oracle.com/cd/E36265_01/index.html

- Oracle Fabric Interconnect F1-15

  http://docs.oracle.com/cd/E38500_01/index.html

- Oracle Integrated Lights Out Manager (ILOM) 3.1

  http://docs.oracle.com/cd/E24707_01/index.html

- Oracle VM

  http://docs.oracle.com/cd/E35328_01/index.html

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Feedback

Provide feedback about this documentation at:

http://www.oracle.com/goto/docfeedback

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Document Revision

Document generated on: 2015-11-27 (revision: 888)

Chapter 1 Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance

Table of Contents

1.1 What is Oracle Private Cloud Appliance 1
1.2 Hardware Components 2
1.2.1 Management Nodes 3
1.2.2 Compute Nodes 4
1.2.3 Storage Appliance 5
1.2.4 Network Infrastructure 6
1.3 Software Components 8
1.3.1 Oracle Private Cloud Appliance Dashboard 8
1.3.2 Password Manager (Wallet) 8
1.3.3 Oracle VM Manager 9
1.3.4 Operating Systems 9
1.3.5 Databases 9
1.3.6 Oracle Private Cloud Appliance Management Software 12
1.3.7 Oracle Private Cloud Appliance Diagnostics Tool 13
1.4 Provisioning and Orchestration 14
1.4.1 Appliance Management Initialization 14
1.4.2 Compute Node Discovery and Provisioning 14
1.4.3 Server Pool Readiness 15
1.5 High Availability 16
1.6 Oracle Private Cloud Appliance Backup 17
1.7 Oracle Private Cloud Appliance Software Update 18

This chapter describes what Oracle Private Cloud Appliance (PCA) is, which hardware and software it consists of, and how it is deployed as a virtualization platform.

1.1 What is Oracle Private Cloud Appliance

Responding to the Cloud Challenges

Cloud architectures and virtualization solutions have become highly sophisticated and complex to implement. They require a skill set that no single administrator has had to master in traditional data centers: system hardware, operating systems, network administration, storage management, applications. Without expertise in every single one of those domains, an administrator cannot take full advantage of the features and benefits of virtualization technology. This often leads to poor implementations with sub-optimal performance and reliability, which impairs the flexibility of a business.

Aside from the risks created by technical complexity and lack of expertise, companies also suffer from an inability to deploy new infrastructure quickly enough to suit their business needs. The administration involved in the deployment of new systems, and the time and effort to configure these systems, can amount to weeks. Provisioning new applications into flexible virtualized environments, in a fraction of the time required for physical deployments, generates substantial financial benefits.

Fast Deployment of Converged Infrastructure

Oracle Private Cloud Appliance is an offering that industry analysts refer to as a *Converged Infrastructure Appliance*: an infrastructure solution in the form of a hardware appliance that comes from the factory

pre-configured. It enables the operation of the entire system as a single unit, not a series of individual servers, network hardware and storage providers. Installation, configuration, high availability, expansion and upgrading are automated and orchestrated as much as possible. Within a few hours after power-on, the appliance is ready to create virtual servers. Virtual servers are commonly deployed from Oracle VM templates (individual pre-configured VMs) and assemblies (interconnected groups of pre-configured VMs).

Modular Implementation of a Complete Stack

With Oracle Private Cloud Appliance (PCA), Oracle offers a unique full stack of hardware, software, virtualization technology and rapid application deployment through assemblies. All this is packaged in a single modular and extensible product. The minimum configuration consists of a base rack with infrastructure components, a pair of management nodes, and two compute nodes. This configuration can be extended by one compute node at a time. All rack units, whether populated or not, are pre-cabled and pre-configured at the factory in order to facilitate the installation of expansion compute nodes on-site at a later time.

Ease of Use

The primary value proposition of Oracle PCA is the integration of components and resources for the purpose of ease of use and rapid deployment. It should be considered a general purpose solution in the sense that it supports the widest variety of operating systems, including Windows, and any application they might host. Customers can attach their existing storage or provide storage solutions from Oracle or third parties.

1.2 Hardware Components

The Oracle Private Cloud Appliance consists of a Sun Rack II 1242 base, populated with the hardware components identified in Figure 1.1.

Note

Oracle Private Cloud Appliance base racks are factory installed with Oracle Server X5-2 nodes. Earlier versions of the Oracle PCA were shipped with Sun Server X4-2 or Sun Server X3-2 nodes, which continue to be supported by the Oracle PCA controller software. Oracle PCA racks that contain Sun Server X3-2 and/or Sun Server X4-2 nodes may also contain Oracle Server X5-2 nodes if additional nodes have been installed, or existing nodes have been replaced.

Software support for Oracle Server X5-2 expansion compute nodes is available as of Release 2.0.2 of the Oracle PCA controller software. Release 2.0.3 adds support for Oracle Server X5-2 management nodes. The software must be upgraded to the correct version to support Oracle Server X5-2 hardware.

Note

The initial version of the Oracle Private Cloud Appliance, the X3-2 base rack, was shipped with the Sun ZFS Storage Appliance 7320. This hardware component has since been replaced with the Oracle ZFS Storage Appliance ZS3-ES. Software support for the Oracle ZFS Storage Appliance ZS3-ES is available as of Release 1.1.3. All software releases continue to support the Sun ZFS Storage Appliance 7320.

Figure 1.1 Components of an Oracle Private Cloud Appliance Rack

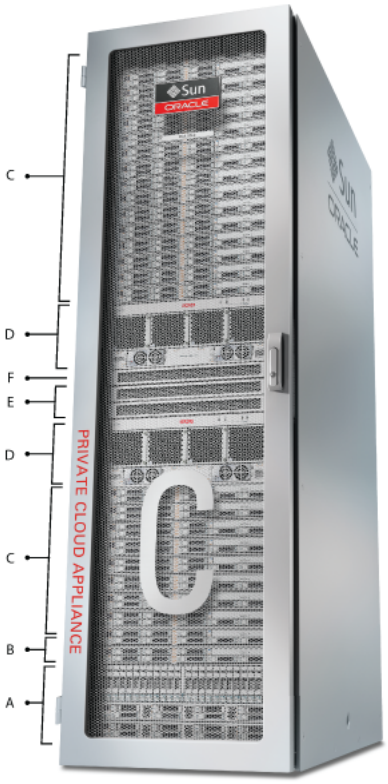

Table 1.1 Figure Legend

Item	Quantity	Description
A	1	Either Oracle ZFS Storage Appliance ZS3-ES or Sun ZFS Storage Appliance 7320
B	2	Either Oracle Server X5-2, Sun Server X4-2, or Sun Server X3-2, used as management nodes
C	2-25	Either Oracle Server X5-2, Sun Server X4-2, or Sun Server X3-2, used as virtualization compute nodes (Due to the higher power requirements of the Oracle Server X5-2, if the appliance is equipped with 22kVA PDUs, the maximum number of compute nodes is 23.)
D	2	Oracle Fabric Interconnect F1-15
E	2	NM2-36P Sun Datacenter InfiniBand Expansion Switch
F	2	Oracle Switch ES1-24

1.2.1 Management Nodes

At the heart of each Oracle PCA installation is a pair of management nodes. They are installed in rack units 5 and 6 and form a cluster in active/standby configuration for high availability: both servers are capable of running the same services and have equal access to the system configuration, but one operates as the master while the other is ready to take over the master functions in case a failure occurs. The master management node runs the full set of services required, while the standby management node runs a

subset of services until it is promoted to the master role. The master role is determined at boot through OCFS2 Distributed Lock Management on an iSCSI LUN, which both management nodes share on the ZFS storage appliance installed at the bottom of the rack. Because rack units are numbered from the bottom up, and the bottom four are occupied by the ZFS Storage Appliance, the master management node is typically the server in rack unit 5. It is the only server that must be powered on by the administrator in the entire process to bring the appliance online.

For details about how high availability is achieved with Oracle PCA, refer to Section 1.5, “High Availability”.

When you power on the Oracle PCA for the first time, you can change the factory default IP configuration of the management node cluster, so that it can be easily reached from your data center network. The management nodes share a Virtual IP, where the management web interface can be accessed. This virtual IP is assigned to whichever server has the *master* role at any given time. During system initialization, after the management cluster is set up successfully, the master management node loads a number of Oracle Linux 6 services, in addition to Oracle VM and its associated MySQL database – including network, sshd, ntpd, iscsi initiator, dhcpd – to orchestrate the provisioning of all system components. During provisioning, all networking and storage is configured, and all compute nodes are discovered, installed and added to an Oracle VM server pool. All provisioning configurations are preloaded at the factory and should not be modified by the customer.

For details about the provisioning process, refer to Section 1.4, “Provisioning and Orchestration”.

1.2.2 Compute Nodes

The compute nodes in the Oracle PCA constitute the virtualization platform. The compute nodes provide the processing power and memory capacity for the virtual servers they host. The entire provisioning process is orchestrated by the management nodes: compute nodes are installed with Oracle VM Server 3.2.9 and additional packages for InfiniBand and Software Defined Networking. When provisioning is complete, the Oracle PCA software expects all compute nodes in the same rack to be part of the same Oracle VM server pool.

For hardware configuration details of the Oracle Server X5-2, Sun Server X4-2 and Sun Server X3-2 compute nodes, refer to Server Components in the Oracle Private Cloud Appliance Installation Guide. Different generations of servers may be mixed within the same installation. In such configurations the version of the Oracle PCA controller software must support the most recent server model installed. Compute nodes of different hardware generations operate within the same server pool but belong to different CPU compatibility groups. Since live migration between CPU compatibility groups is not supported, virtual machines have to be cold-migrated between compute nodes of different generations. For more information about CPU compatibility groups, please refer to the section “Server Processor Compatibility Groups” in the Oracle VM User's Guide.

The Oracle PCA Dashboard allows the administrator to monitor the health and status of the compute nodes, as well as all other rack components, and perform certain system operations. The virtual infrastructure is configured and managed with Oracle VM Manager.

The Oracle PCA offers modular compute capacity that can be increased according to business needs. The minimum configuration of the base rack contains just two compute nodes, but it can be expanded by one node at a time up to 25 compute nodes. Apart from the hardware installation, adding compute nodes requires no intervention by the administrator. New nodes are discovered, powered on, installed and provisioned automatically by the master management node. The additional compute nodes are integrated into the existing configuration and, as a result, the Oracle VM server pool offers increased capacity for more or larger virtual machines.

Because it is difficult to quantify the compute capacity as a number of virtual machines, it is suggested that you base calculations on the amount of RAM as a rule of thumb. With 256GB per compute node and

a minimal *dom0* overhead, it is safe to assume that each compute node can host 60 virtual machines with 4GB RAM, or 30 virtual machines with 8GB RAM, or 15 virtual machines with 16GB RAM, or any combination that adds up to an amount of RAM just a couple of GB under physical memory of the server.

1.2.3 Storage Appliance

The Oracle ZFS Storage Appliance ZS3-ES installed at the bottom of the appliance rack should be considered a 'system disk' for the entire appliance. Its main purpose is to provide storage space for the Oracle PCA software. A portion of the disk space is made available for customer use and is sufficient for an Oracle VM storage repository with a limited number of virtual machines, templates and assemblies.

The hardware configuration of the Oracle ZFS Storage Appliance ZS3-ES is as follows:

- Two clustered storage heads with two 1.6TB SSDs each, used exclusively for cache and logging
- One fully populated disk chassis with twenty 900GB SATA hard disks
- RAID-Z2 configuration, for best balance between performance and data protection, with a total usable space of 11.3TB

Note

Oracle PCA Release 1.0 base racks, which may be upgraded to the Release 2.0 software stack, use a Sun ZFS Storage Appliance 7320. It offers the same performance, functionality, and configuration, but its storage heads use smaller SSDs. The disk shelf and its disks are identical in both models.

The storage appliance is connected to the management subnet (`192.168.4.0/24`) and the InfiniBand (IPoIB) storage subnet (`192.168.40.0/24`). Because both heads form a cluster, they share a single IP in each subnet. The RAID-Z2 storage pool contains two projects, named `OVCA` and `OVM` .

The `OVCA` project contains all LUNs and file systems used by the Oracle PCA software:

- LUNs
  - `Locks` (12GB) – to be used exclusively for cluster locking on the two management nodes
  - `Manager` (200GB) – to be used exclusively as an additional file system on both management nodes
- File systems:
  - `MGMT_ROOT` – to be used for storage of all files specific to the Oracle PCA
  - `Database` – to be used for all system databases
  - `Incoming` (20GB) – to be used for FTP file transfers, primarily for Oracle PCA component backups
  - `Templates` – placeholder file system for future use
  - `User` – placeholder file system for future use
  - `Yum` – placeholder file system for future use

The `OVM` project contains all LUNs and file systems used by Oracle VM:

- LUNs
  - `iscsi_repository1` (300GB) – to be used as Oracle VM storage repository

- `iscsi_serverpool1` (12GB) – to be used as server pool file system for the Oracle VM clustered server pool

- File systems:
  - `nfs_repository1` (300GB) – to be used as Oracle VM storage repository in case NFS is preferred over iSCSI
  - `nfs_serverpool1` (12GB) – to be used as server pool file system for the Oracle VM clustered server pool in case NFS is preferred over iSCSI

In addition to offering storage, the ZFS storage appliance also runs the xinetd and tftpd services. These complement the Oracle Linux services on the master management node in order to orchestrate the provisioning of all Oracle PCA system components.

1.2.4 Network Infrastructure

The Oracle Private Cloud Appliance relies on a combination of Ethernet connectivity and an InfiniBand network fabric. The appliance rack contains redundant network hardware components, which are pre-cabled at the factory to help ensure continuity of service in case a failure should occur.

Ethernet

The Ethernet network relies on two interconnected Oracle Switch ES1-24 switches, to which all other rack components are connected with CAT6 Ethernet cables. This network serves as the appliance management network, in which every component has a predefined IP address in the `192.168.4.0/24` range. In addition, all management and compute nodes have a second IP address in this range, which is used for Oracle Integrated Lights Out Manager (ILOM) connectivity.

While the appliance is initializing, the InfiniBand fabric is not accessible, which means that the management network is the only way to connect to the system. Therefore, the administrator should connect a workstation to the available Ethernet port 19 in one of the Oracle Switch ES1-24 switches, and assign the fixed IP address `192.168.4.254` to the workstation. From this workstation, the administrator opens a browser connection to the web server on the master management node at `http://192.168.4.216` , in order to monitor the initialization process and perform the initial configuration steps when the appliance is powered on for the first time.

InfiniBand

The Oracle PCA rack contains two NM2-36P Sun Datacenter InfiniBand Expansion Switches. These redundant switches have redundant cable connections to both InfiniBand ports in each management node, compute node and storage head. Both InfiniBand switches, in turn, have redundant cable connections to both Fabric Interconnects in the rack. All these components combine to form a physical InfiniBand backplane with a 40Gbit (Quad Data Rate) bandwidth.

When the appliance initialization is complete, all necessary Oracle PCA software packages, including host drivers and InfiniBand kernel modules, have been installed and configured on each component. At this point, the system is capable of using software defined networking (SDN) configured on top of the physical InfiniBand fabric. SDN is implemented through the Fabric Interconnects.

Fabric Interconnect

All Oracle PCA network connectivity is managed through the Fabric Interconnects. Data is transferred across the physical InfiniBand fabric, but connectivity is implemented in the form of Software Defined Networks (SDN), which are sometimes referred to as 'clouds'. The physical InfiniBand backplane is

capable of hosting thousands of virtual networks. These Private Virtual Interconnects (PVI) dynamically connect virtual machines and bare metal servers to networks, storage and other virtual machines, while maintaining the traffic separation of hard-wired connections and surpassing their performance.

During the initialization process of the Oracle PCA, five essential SDNs are configured: a storage network, an Oracle VM management network, a management Ethernet network, and two VLAN-enabled virtual machine networks.

- The **storage network** is a bonded IPoIB connection between the management nodes and the ZFS storage appliance, and uses the `192.168.40.0/24` subnet. This network also fulfills the heartbeat function for the clustered Oracle VM server pool. DHCP ensures that compute nodes are assigned an IP address in this subnet.
- The **Oracle VM management network** is a PVI that connects the management nodes and compute nodes in the `192.168.140.0/24` subnet. It is used for all network traffic inherent to Oracle VM Manager, Oracle VM Server and the Oracle VM Agents.
- The **management Ethernet network** is a bonded Ethernet connection between the management nodes. The primary function of this network is to provide access to the management nodes from the data center network, and enable the management nodes to run a number of system services. Since all compute nodes are also connected to this network, Oracle VM can use it for virtual machine connectivity, with access to and from the data center network. The management node external network settings are configurable through the Network Setup tab in the Oracle PCA Dashboard. If this network is a VLAN, its ID or tag must be configured in the Network Setup tab of the Dashboard.
- The **public virtual machine network** is a bonded Ethernet connection between the compute nodes. Oracle VM uses this network for virtual machine connectivity, where external access is required. VLAN 1 is automatically configured for this network. Customers can add their own VLANs to the Oracle VM network configuration, and define the subnet(s) appropriate for IP address assignment at the virtual machine level. For external connectivity, the next-level data center switches must be configured to accept your tagged VLAN traffic.
- The **private virtual machine network** is a bonded Ethernet connection between the compute nodes. Oracle VM uses this network for virtual machine connectivity, where only internal access is required. VLAN 1 is automatically configured for this network. Customers can add VLANs of their choice to the Oracle VM network configuration, and define the subnet(s) appropriate for IP address assignment at the virtual machine level.

Finally, the Fabric Interconnects also manage the physical public network connectivity of the Oracle PCA. Two 10GbE ports on each Fabric Interconnect must be connected to redundant next-level data center switches. At the end of the initialization process, the administrator assigns three reserved IP addresses from the data center (public) network range to the management node cluster of the Oracle PCA: one for each management node, and an additional Virtual IP shared by the clustered nodes. From this point forward, the Virtual IP is used to connect to the master management node's web server, which hosts both the Oracle PCA Dashboard and the Oracle VM Manager web interface.

Caution

It is critical that **both** Fabric Interconnects have **two** 10GbE connections each to a pair of next-level data center switches. This configuration with four cable connections provides redundancy and load splitting at the level of the Fabric Interconnects, the 10GbE ports and the data center switches. This outbound cabling should not be crossed or meshed, because the internal connections to the pair of Fabric Interconnects are already configured that way. The cabling pattern plays a key role in the continuation of service during failover scenarios involving Fabric Interconnect outages and other components.

1.3 Software Components

This section describes the main software components the Oracle Private Cloud Appliance uses for operation and configuration.

1.3.1 Oracle Private Cloud Appliance Dashboard

The Oracle PCA provides its own web-based graphical user interface that can be used to perform a variety of administrative tasks specific to the appliance. The Oracle PCA Dashboard is a WebLogic application that is available via the active management node. It is installed on top of the Oracle WebLogic Server 11*g* that is packaged with Oracle VM Manager, as described in Section 1.3.3, “Oracle VM Manager”.

Use the Dashboard to perform the following tasks:

- Appliance system monitoring and component identification
- Monitoring and identifying physical network connections
- Initial configuration of management node networking data
- Resetting of the global password for Oracle PCA configuration components

The Oracle PCA Dashboard is described in detail in Chapter 2, *Monitoring and Managing Oracle Private Cloud Appliance*.

1.3.2 Password Manager (Wallet)

All components of the Oracle Private Cloud Appliance have administrator accounts with a default password. After applying your data center network settings through the Oracle PCA Dashboard, it is recommended that you modify the default appliance password. The Authentication tab allows you to set a new password, which is applied to the main system configuration components.

Caution

The Oracle WebLogic Server password cannot be set from the Dashboard. Please change the password in the Oracle WebLogic Server Administration Console. Make sure it matches the password configured in the Wallet. Otherwise the Oracle PCA will not be able to connect to Oracle WebLogic Server. For details, see Section 7.9, “Changing Oracle WebLogic Server Passwords”.

Caution

When modifying component ILOM passwords, make sure that it is identical for all compute nodes and management nodes. The server ILOM password must be updated in three locations:

- the ILOM's own interface
- the Oracle PCA Wallet
- the server IPMI configuration in Oracle VM Manager

Passwords for all accounts on all components are stored in a global Wallet, secured with 512-bit encryption. The password update functionality in the Dashboard is restricted to the Dashboard UI itself, Oracle VM Manager, the root user account on both management nodes, and the `ovs` user account of the

Oracle VM mySQL database. It is possible to change the password for other accounts and components by setting the password manually and updating the Wallet entry from the command line of the master management node. For details, see Section 7.10, “Replacing Default Passwords Manually”.

Manually changing passwords is required in case the Wallet is out of sync with the current status of the hardware. For instance, this could occur after a backup of the appliance software and configuration is restored.

1.3.3 Oracle VM Manager

All virtual machine management tasks are performed within Oracle VM Manager, a WebLogic application that is installed on each of the management nodes and which provides a web-based management user interface and a command line interface that allows you to manage your Oracle VM infrastructure within the Oracle PCA.

Oracle VM Manager is comprised of the following software components:

- **Oracle VM Manager application**: provided as an Oracle WebLogic Server domain and container.
- **Oracle WebLogic Server 11g**: including Application Development Framework (ADF) Release 11*g*, used to host and run the Oracle VM Manager application
- **MySQL 5.5 Enterprise Edition Server**: for the exclusive use of the Oracle VM Manager application as a management repository and installed on the Database file system hosted on the ZFS storage appliance.

Administration of virtual machines is performed using the Oracle VM Manager web user interface, as described in Chapter 4, *Managing the Oracle VM Virtual Infrastructure*. While it is possible to use the command line interface provided with Oracle VM Manager, this is considered an advanced activity that should only be performed with a thorough understanding of the limitations of Oracle VM Manager running in the context of an Oracle Private Cloud Appliance.

1.3.4 Operating Systems

Hardware components of the Oracle Private Cloud Appliance run their own operating systems:

- Management Nodes: Oracle Linux 6 with UEK2
- Compute Nodes: Oracle Linux 5 with UEK2
- Oracle ZFS Storage Appliance ZS3-ES: Oracle Solaris 11

All other components run a particular revision of their respective firmware. All operating software has been selected and developed to work together as part of the Oracle PCA. When an update is released, the appropriate versions of all software components are bundled. When a new software release is activated, all component operating software is updated accordingly. You should not attempt to update individual components unless Oracle explicitly instructs you to.

1.3.5 Databases

The Oracle PCA uses a number of databases to track system states, handle configuration and provisioning, and for Oracle VM Manager. All databases are stored on the ZFS storage appliance, and are exported via an NFS file system. The databases are accessible to each management node to ensure high availability.

The following table lists the different databases used by the Oracle PCA.

Table 1.2 Oracle Private Cloud Appliance Databases

Item	Description
Oracle PCA Node Database	Contains information on every compute node and management node in the rack, including the state used to drive the provisioning of compute nodes and data required to handle software updates. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/node` on the ZFS, accessible via `/nfs/shared_storage/db/node` on each management node
Oracle Fabric Interconnect Database	Contains IP and host name data for the Oracle Fabric Interconnect F1-15s. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/infrastructure` on the ZFS, accessible via `/nfs/shared_storage/db/infrastructure` on each management node
Oracle PCA Netbundle Database	Predefines Ethernet and bond device names for all possible networks that can be configured throughout the system, and which are allocated dynamically. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/netbundle` on the ZFS, accessible via `/nfs/shared_storage/db/netbundle` on each management node
Oracle Switch ES1-24 Ports Database	Defines the factory-configured map of Oracle Switch ES1-24 ports to the rack unit or element to which that port is connected. It is used to map Oracle Switch ES1-24 ports to machine names. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/opus_ports` on the ZFS, accessible via `/nfs/shared_storage/db/opus_ports` on each management node
Oracle PCA DHCP Database	Contains information on the assignment of DHCP addresses to newly detected compute nodes. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/dhcp` on the ZFS, accessible via `/nfs/shared_storage/db/dhcp` on each management node
Oracle PCA Mini Database	A multi-purpose database used to map compute node hardware profiles to on-board disk size information. It also contains valid hardware configurations that servers must comply with in order to be accepted as an Oracle PCA component. Entries contain a sync ID for more convenient usage within the CLI. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/mini_db` on the ZFS, accessible via `/nfs/shared_storage/db/mini_db` on each management node

Item	Description
Oracle PCA Setup Database	Contains the data set by the Oracle PCA Dashboard setup facility. The data in this database is automatically applied by both the active and standby management nodes when a change is detected. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/setup` on the ZFS, accessible via `/nfs/shared_storage/db/setup` on each management node
Oracle PCA Task Database	Contains state data for all of the asynchronous tasks that have been dispatched within the Oracle PCA. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/task` on the ZFS, accessible via `/nfs/shared_storage/db/task` on each management node
Oracle PCA Synchronization Databases	Contain data and configuration settings for the synchronization service to apply and maintain across rack components. Errors from failed attempts to synchronize configuration parameters across appliance components can be reviewed in the `sync_errored_tasks` database, from where they can be retried or acknowledged. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/sync_*` on the ZFS, accessible via `/nfs/shared_storage/db/sync_*` on each management node
Oracle PCA Update Database	Used to track the two-node coordinated management node update process. **Note** Database schema changes and wallet changes between different releases of the controller software are written to a file. It ensures that these critical changes are applied early in the software update process, before any other appliance components are brought back up. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/update` on the ZFS, accessible via `/nfs/shared_storage/db/update` on each management node
Oracle VM Manager Database	Used on each management node as the management database for Oracle VM Manager. It contains all configuration details of the Oracle VM environment (including servers, pools, storage and networking), as well as the virtualized systems hosted by the environment. **Type:** MySQL Database **Location:** `MGMT_ROOT/ovmm_mysql/data/` on the ZFS, accessible via `/nfs/shared_storage/ovmm_mysql/data/` on each management node

1.3.6 Oracle Private Cloud Appliance Management Software

The Oracle Private Cloud Appliance includes software that is designed for the provisioning, management and maintenance of all of the components within the appliance.

Important

This software is not designed for human interaction. All configuration and management tasks must be performed using the Oracle PCA Dashboard. Do not attempt to run any of these processes directly without explicit instruction from an Oracle Support representative. Attempting to do so may render your appliance unusable.

This software largely consists of a number of Python applications that run on the active management node. These applications are found in `/usr/sbin` on each management node and are listed as follows:

- `pca-backup`: the script responsible for performing backups of the appliance configuration as described in Section 1.6, “Oracle Private Cloud Appliance Backup”
- `pca-check-master`: a script that verifies which of the two management nodes currently has the master role
- `ovca-daemon`: the core provisioning and management daemon for the Oracle PCA
- `pca-dhcpd`: a helper script to assist the DHCP daemon with the registration of compute nodes
- `pca-diag`: a tool to collect diagnostic information from your Oracle PCA, as described in Section 1.3.7, “Oracle Private Cloud Appliance Diagnostics Tool”
- `pca-factory-init`: the appliance initialization script used to set the appliance to its factory configuration
- `pca-helper`: a support script used to perform particular configuration tasks and to assist with appliance configuration
- `pca-node-db`: a script used to maintain and update the Oracle PCA Node Database described in Section 1.3.5, “Databases”
- `pca-redirect`: a script to redirect HTTP or HTTPS requests to the Oracle PCA Dashboard described in Section 1.3.1, “Oracle Private Cloud Appliance Dashboard”
- `ovca-remote-rpc`: a script that allows other components to communicate with the various management scripts available on the management node
- `ovca-rpc`: a script that allows the Oracle PCA Dashboard to communicate directly with the underlying management scripts running on the management node
- `pca-update-password`: a script that controls password updates for each different component in the appliance
- `pca-updater`: the script that allows you to download updated software for the Oracle PCA. It also allows you to activate the update, as described in Section 1.7, “Oracle Private Cloud Appliance Software Update”
- `pca-vm-tool`: an internal script to facilitate virtual machine deployment.

Many of these applications use a specific Oracle PCA library that is installed in `/usr/lib/python2.6/site-packages/ovca/` on each management node.

1.3.7 Oracle Private Cloud Appliance Diagnostics Tool

The Oracle Private Cloud Appliance includes a tool that can be run to collect diagnostic data: logs and other types of files that can help to troubleshoot hardware and software problems. This tool is located in `/usr/sbin/` on each management and compute node, and is named `pca-diag`. The data it retrieves, depends on the selected command line arguments:

- `pca-diag`

  When you enter this command, without any additional arguments, the tool retrieves a basic set of files that provide insights into the current health status of the Oracle PCA. You can run this command on all management and compute nodes. All collected data is stored in `/tmp`, compressed into a single tarball (`pcadiag_<node-hostname>_<ID>_<date>_<time>.tar.bz2`).

- `pca-diag version`

  When you enter this command, version information for the current Oracle PCA software stack is displayed. The `version` argument cannot be combined with any other argument.

- `pca-diag ilom`

  When you enter this command, diagnostic data is retrieved, by means of `ipmitool`, through the ILOM of each component in the rack. The data set includes details about the host's operating system, processes, health status, hardware and software configuration, as well as a number of files specific to the Oracle PCA configuration. You can run this command on all management and compute nodes. All collected data is stored in `/tmp`, compressed into a single tarball (`pcadiag_<node-hostname>_<ID>_<date>_<time>.tar.bz2`).

- `pca-diag vmpinfo`

  Caution

  When using the `vmpinfo` argument, the command must be run from the master management node.

  When you enter this command, the Oracle VM diagnostic data collection mechanism is activated. The `vmpinfo3` script collects logs and configuration details from the Oracle VM Manager, and logs and `sosreport` information from each Oracle VM Server or compute node. All collected data is stored in `/tmp`, compressed into two tarballs: `pcadiag_<node-hostname>_<ID>_<date>_<time>.tar.bz2` and `vmpinfo3-<version>-<date>-<time>.tar.gz`.

  To collect diagnostic information for a subset of the Oracle VM Servers in the environment, you run the command with an additional `servers` parameter: `pca-diag vmpinfo servers='ovcacn07r1,ovcacn08r1,ovcacn09r1'`

Diagnostic collection with `pca-diag` is possible from the command line of any node in the system. Only the master management node allows you to use all of the command line arguments. Although `vmpinfo` is not available on the compute nodes, running `pca-diag` directly on the compute can help retrieve important diagnostic information regarding Oracle VM Server that cannot be captured with `vmpinfo`. Running `pca-diag` from different locations can be particularly helpful in debugging InfiniBand fabric issues.

The `pca-diag` tool is typically run by multiple users with different roles. System administrators or field service engineers may use it as part of their standard operating procedures, or Oracle Support teams may request that the tool be run in a specific manner as part of an effort to diagnose and resolve reported hardware or software issues. For additional information and instructions, also refer to the section "Data Collection for Service and Support" in the Oracle Private Cloud Appliance Release Notes.

1.4 Provisioning and Orchestration

As a converged infrastructure solution, the Oracle Private Cloud Appliance is built to eliminate many of the intricacies of optimizing the system configuration. Hardware components are installed and cabled at the factory. Configuration settings and installation software are preloaded onto the system. Once the appliance is connected to the data center power source and public network, the provisioning process between the administrator pressing the power button of the first management node and the appliance reaching its *Deployment Readiness* state is entirely orchestrated by the master management node. This section explains what happens as the Oracle PCA is initialized and all nodes are provisioned.

1.4.1 Appliance Management Initialization

Boot Sequence and Health Checks

When power is applied to the first management node, it takes approximately five minutes for the server to boot. While the Oracle Linux 6 operating system is loading, an Apache web server is started, which serves a static welcome page the administrator can browse to from the workstation connected to the appliance management network.

The necessary Oracle Linux services are started as the server comes up to runlevel 3 (multi-user mode with networking). At this point, the management node executes a series of system health checks. It verifies that all expected infrastructure components are present on the appliance management network and in the correct predefined location, identified by the rack unit number and fixed IP address. Next, the management node probes the ZFS storage appliance for a management NFS export and a management iSCSI LUN with OCFS2 file system. The storage and its access groups have been configured at the factory. If the health checks reveal no problems, the `ocfs2` and `o2cb` services are started up automatically.

Management Cluster Setup

When the OCFS2 file system on the shared iSCSI LUN is ready, and the `o2cb` services have started successfully, the management nodes can join the cluster. In the meantime, the first management node has also started the second management node, which will come up with an identical configuration. Both management nodes eventually join the cluster, but the first management node will take an exclusive lock on the shared OCFS2 file system using Distributed Lock Management (DLM). The second management node remains in permanent standby and takes over the lock only in case the first management node goes down or otherwise releases its lock.

With mutual exclusion established between both members of the management cluster, the master management node continues to load the remaining Oracle PCA services, including dhcpd, Oracle VM Manager and the Oracle PCA databases. The virtual IP address of the management cluster is also brought online, and the Oracle PCA Dashboard is started within WebLogic. The static Apache web server now redirects to the Dashboard at the virtual IP, where the administrator can access a live view of the appliance rack component status.

Once the `dhcpd` service is started, the system state changes to *Provision Readiness*, which means it is ready to discover non-infrastructure components.

1.4.2 Compute Node Discovery and Provisioning

Node Manager

To discover compute nodes, the Node Manager on the master management node uses a DHCP server and the node database. The node database is a BerkeleyDB type database, located on the management NFS share, containing the state and configuration details of each node in the system, including MAC

addresses, IP addresses and host names. The discovery process of a node begins with a DHCP request from the ILOM. Most discovery and provisioning actions are synchronous and occur sequentially, while time consuming installation and configuration processes are launched in parallel and asynchronously. The DHCP server hands out pre-assigned IP addresses on the appliance management network (`192.168.4.0/24`). When the Node Manager has verified that a node has a valid service tag for use with Oracle PCA, it launches a series of provisioning tasks. All required software resources have been loaded onto the ZFS storage appliance at the factory.

Provisioning Tasks

The provisioning process is tracked in the node database by means of status changes. The next provisioning task can only be started if the node status indicates that the previous task has completed successfully. For each valid node, the Node Manager begins by building a PXE configuration and forces the node to boot using Oracle PCA runtime services. After the hardware RAID-1 configuration is applied, the node is restarted to perform a kickstart installation of Oracle VM Server. Crucial kernel modules and host drivers for InfiniBand and IO Director (Fabric Interconnect) support are added to the installation. At the end of the installation process, the network configuration files are updated to allow all necessary network interfaces and bonds to be brought up.

Now that the PVI for the Oracle VM management network exists, the compute node is rebooted one last time to reconfigure the Oracle VM Agent to communicate over the PVI. At this point, the node is ready for Oracle VM Manager discovery.

1.4.3 Server Pool Readiness

Oracle VM Server Pool

When the Node Manager detects a fully installed compute node that is ready to join the Oracle VM environment, it issues the necessary Oracle VM CLI commands to add the new node to the Oracle VM server pool. With the discovery of the first node, the system also configures the clustered Oracle VM server pool with the appropriate networking, access to the shared storage, and a virtual IP. For every compute node added to Oracle VM Manager the IPMI configuration is stored in order to enable convenient remote power-on/off.

Oracle PCA expects that all compute nodes in one rack belong to a single clustered server pool with High Availability (HA) and Distributed Resource Scheduling (DRS) enabled. When all compute nodes have joined the Oracle VM server pool, the appliance is in *Ready* state, meaning virtual machines (VMs) can be deployed.

Expansion Compute Nodes

When an expansion compute node is installed, its presence is detected based on the DHCP request from its ILOM. If the new server is identified as an Oracle PCA node, an entry is added in the node database with "*new*" state. This triggers the initialization and provisioning process. New compute nodes are integrated seamlessly to expand the capacity of the running system, without the need for manual reconfiguration by an administrator.

Synchronization Service

As part of the provisioning process, a number of configuration settings are applied, either globally or at individual component level. Some are visible to the administrator, and some are entirely internal to the system. Throughout the life cycle of the appliance, software updates, capacity extensions and configuration changes will occur at different points in time. For example, an expansion compute node may have different hardware, firmware and software compared to the servers already in use in the environment, and it comes with factory default settings that do not match those of the running system. A synchronization

service, implemented on the management nodes, can set and maintain configurable parameters across heterogeneous sets of components within an Oracle PCA environment. It facilitates the integration of new system components in case of capacity expansion or servicing, and allows the administrator to streamline the process when manual intervention is required. The CLI provides an interface to the exposed functionality of the synchronization service.

1.5 High Availability

The Oracle Private Cloud Appliance is designed for high availability at every level of its component make-up.

Management Node Failover

During the factory installation of an Oracle PCA, the management nodes are configured as a cluster. The cluster relies on an OCFS2 file system exported as a LUN from the ZFS storage to perform the heartbeat function and to store a lock file that each management node attempts to take control of. The management node that has control over the lock file automatically becomes the master or active node in the cluster.

When the Oracle PCA is first initialized, the `o2cb` service is started on each management node. This service is the default cluster stack for the OCFS2 file system. It includes a node manager that keeps track of the nodes in the cluster, a heartbeat agent to detect live nodes, a network agent for intra-cluster node communication and a distributed lock manager to keep track of lock resources. All these components are in-kernel.

Additionally, the `ovca` service is started on each management node. The management node that obtains control over the cluster lock and is thereby promoted to the master or active management node, runs the full complement of Oracle PCA services. This process also configures the Virtual IP, that is used to access the active management node, so that it is 'up' on the active management node and 'down' on the standby management node. This ensures that, when attempting to connect to the Virtual IP address that you configured for the management nodes, you are always accessing the active management node.

In the case where the active management node fails, the cluster detects the failure and the lock is released. Since the standby management node is constantly polling for control over the lock file, it detects when it has control of this file and the `ovca` service brings up all of the required Oracle PCA services. The Virtual IP on the standby management node is brought up as it is promoted to the active role.

When the management node that failed comes back online, it no longer has control of the cluster lock file. It is automatically put into standby mode, and the Virtual IP is configured as 'down'. This means that one of the two management nodes in the rack is always available via the same IP address and is always correctly configured. The management node failover process takes up to 5 minutes to complete.

Oracle VM Management Database Failover

The Oracle VM Manager database files are located on a shared file system exposed by the ZFS storage appliance. The active management node runs the MySQL database server, which accesses the database files on the shared storage. In the event that the management node fails, the standby management node is promoted and the MySQL database server on the promoted node is started so that the service can resume as normal. The database contents are available to the newly running MySQL database server.

Compute Node Failover

High availability (HA) of compute nodes within the Oracle PCA is enabled through the clustered server pool that is created automatically in Oracle VM Manager during the compute node provisioning process. Since the server pool is configured as a cluster using an underlying OCFS2 file system, HA-enabled virtual

machines running on any compute node can be migrated and restarted automatically on an alternate compute node in the event of failure.

Storage Redundancy

Further redundancy is provided through the use of the ZFS storage appliance to host storage. This component is configured to support RAID-Z2 providing integrated redundancy with a fault tolerance of up to two failed drives with zero data loss. Furthermore, the storage appliance includes two storage heads or controllers that are interconnected in a clustered configuration. The pair of controllers operate in an active-passive configuration, meaning that in the event the active controller fails, the standby controller can take over all traffic without an interruption in service.

Network Redundancy

All of the customer-usable networking within the Oracle PCA is configured for redundancy. Only the internal administrative Ethernet network, which is used for initialization and ILOM connectivity, is not redundant. There are two of each switch type to ensure that there is no single point of failure. Networking cabling is equally duplicated and switches are interconnected as described in Section 1.2.4, “Network Infrastructure”.

1.6 Oracle Private Cloud Appliance Backup

The configuration of all components within Oracle Private Cloud Appliance is automatically backed up and stored on the ZFS storage appliance as a set of archives. Backups are named with a time stamp for when the backup is run.

During initialization, a crontab entry is created on each management node to perform a global backup twice in every 24 hours. The first backup runs at 09h00 and the second at 21h00. Only the active management node actually runs the backup process when it is triggered.

Note

To trigger a backup outside of the default schedule, use the Command Line Interface. For details, refer to Section 3.2.1, “backup”.

Backups are stored on the `MGMT_ROOT` file system on the ZFS storage appliance and are accessible on each management node at `/nfs/shared_storage/backups`. When the backup process is triggered, it creates a directory named with the time stamp for the current backup process. Within this directory several subdirectories are also created:

- **nm2**: contains the NM2-36P Sun Datacenter InfiniBand Expansion Switch configuration data
- **opus**: contains the Oracle Switch ES1-24 configuration data
- **ovca**: contains all of the configuration information relevant to the deployment of the management nodes such as the password wallet, the network configuration of the management nodes, configuration databases for the Oracle PCA services, and DHCP configuration.
- **ovmm**: contains the most recent backup of the Oracle VM Manager database, the actual source data files for the current database, and the UUID information for the Oracle VM Manager installation. Note that the actual backup process for the Oracle VM Manager database is handled automatically from within Oracle VM Manager and is described in detail in the section entitled Oracle VM Manager MySQL Backup, in the Oracle VM Installation and Upgrade Guide.
- **xsigo**: contains the configuration data for the Fabric Interconnects.

- **zfssa**: contains all of the configuration information for the ZFS storage appliance

The backup process collects data for each component in the appliance and ensures that it is stored in a way that makes it easy to restore that component to operation in the case of failure [1] .

Taking regular backups is standard operating procedure for any production system. The internal backup mechanism cannot protect against full system failure, site outage or disaster. Therefore, you should consider implementing a backup strategy to copy key system data to external storage. This requires a machine with connections to the internal appliance networks as well as the external (public) data center network.

For a detailed description of the backup contents, and for guidelines to export internal backups outside the appliance, refer to the Oracle technical white paper entitled *Oracle Private Cloud Appliance Backup and Recovery Guide*.

1.7 Oracle Private Cloud Appliance Software Update

Oracle Private Cloud Appliance includes the tools required to perform a full software update of the appliance. On the master management node command line, commands are issued to download a new version of the software stack from a specified URI and update all components to the latest version. For step-by-step instructions, see Section 2.6, “Update”.

An Oracle PCA software stack image is provided in the form of a zip file, approximately 4.6 GB in size. The zip file contains an automatically built ISO with the latest software updates from Oracle.

The software image can be downloaded using the FTP, HTTP or HTTPS protocols [2] . Once the image has been downloaded it is extracted into the `mgmt_image`, `hmp_image` and `ovm_image` subdirectories of `/nfs/shared_storage`, which resides in the `MGMT_ROOT` file system on the ZFS storage appliance.

From the command line, you are able to *activate* a downloaded software image. Activation starts a coordinated process on the management nodes that relies on management node failover to succeed.

The activation process can be summarized in the following way:

- The current master begins the process by creating a database entry to include a message containing the version number of the software that you are updating to; a time stamp and the node name of the inactive management node.
- The current master reboots the inactive management node and provisions it with the new software image in the same way that a compute node is provisioned at start-up. The current software on the inactive management node is wiped and the new software is installed directly from the ISO stored in `/nfs/shared_storage/mgmt_image` on the storage appliance.
- After the install, the management node is rebooted and starts its *ovca* service. This causes it to re-enter the remaster code and it finalizes its own database entry.
- The current master detects the database update indicating that the installation has completed and the newly installed management node is ready to resume service as the new master. It creates its own database entry to indicate that it is now being updated, and then reboots.
- The newly installed management node is promoted to the cluster master and then proceeds to perform the provisioning required to install the new software image on the rebooting management node.

[1] Restoration from backup must only be performed by Oracle Service Personnel.
[2] It may be necessary to edit the Oracle PCA system properties, using the CLI on each management node, to ensure that the correct proxy settings are specified for a download to succeed. See Section 7.1, “Adding Proxy Settings for Oracle Private Cloud Appliance Updates” for more information.

- When installation is complete, the management node is rebooted and comes up in standby mode.

Chapter 2 Monitoring and Managing Oracle Private Cloud Appliance

Table of Contents

2.1 Connecting and Logging in to the Oracle Private Cloud Appliance Dashboard 22
2.2 Hardware View 22
2.3 Network View 25
2.4 Network Settings 25
2.5 Authentication 27
2.6 Update 28
2.6.1 Mandatory Updates 29
2.6.2 Updating the Oracle PCA Controller Software 30
2.6.3 Upgrading the Virtualization Platform 33
2.6.4 Upgrading Component Firmware 35

Monitoring and management of the Oracle Private Cloud Appliance (PCA) is achieved using the Oracle PCA Dashboard. This web-based graphical user interface is also used to perform the initial configuration of the appliance beyond the instructions provided in the Quick Start poster included in the packaging of the appliance.

Warning

Before starting the system and applying the initial configuration, read and understand the Oracle Private Cloud Appliance Release Notes. The section Known Limitations and Workarounds provides information that is critical for correctly executing the procedures in this document. Ignoring the release notes may cause you to configure the system incorrectly. Bringing the system back to normal operation may require a complete factory reset.

The Oracle PCA Dashboard allows you to perform the following tasks:

- Initial software configuration (and reconfiguration) for the appliance using the Network Setup tab, as described in Section 2.4, “Network Settings”.
- Hardware monitoring and identification of each hardware component used in the appliance, accessed via the Hardware View tab described in Section 2.2, “Hardware View”.
- Physical network monitoring and identification via the Network View tab described in Section 2.3, “Network View”.
- Reconfiguration of appliance software, such as resetting the global password used for different components within the appliance, via the Authentication tab, as described in Section 2.5, “Authentication”.

The Oracle PCA software includes functionality that is currently not available through the Dashboard user interface:

- **Backup**

  The configuration of all components within Oracle PCA is automatically backed up based on a crontab entry. This functionality is not configurable. Restoring a backup requires the intervention of an Oracle-qualified service person. For details, see Section 1.6, “Oracle Private Cloud Appliance Backup”.

- **Update**

  The update process is controlled from the command line of the master management node. For details, see Section 1.7, “Oracle Private Cloud Appliance Software Update”. For step-by-step instructions, see Section 2.6, “Update”.

2.1 Connecting and Logging in to the Oracle Private Cloud Appliance Dashboard

To open the Login page of the Oracle PCA Dashboard, enter the following address in a Web browser:

https://*manager-vip*:7002/dashboard

Where, *manager-vip* refers to the shared Virtual IP address that you have configured for your management nodes during installation. By using the shared Virtual IP address, you ensure that you always access the Oracle PCA Dashboard on the active management node.

Note

If you are following the installation process and this is your first time accessing the Oracle PCA Dashboard, the Virtual IP address in use by the master management node is set to the factory default `192.168.4.216` . This is an IP address in the internal appliance management network, which can only be reached if you use a workstation patched directly into the available Ethernet port 19 in one of the Oracle Switch ES1-24 switches. The default user name is **admin** and the default password is **Welcome1**.

Important

You must ensure that if you are accessing the Oracle PCA Dashboard through a firewalled connection, the firewall is configured to allow TCP traffic on the port that the Oracle PCA Dashboard is using to listen for connections.

Enter your Oracle PCA Dashboard administration user name in the **Username** field. This is the administration user name you configured during installation. Enter the password for the Oracle PCA Dashboard administration user name in the **Password** field.

Important

The Oracle PCA Dashboard makes use of cookies in order to store session data. Therefore, to successfully login and use the Oracle PCA Dashboard your web browser must accept cookies from the Oracle PCA Dashboard host.

2.2 Hardware View

The **Hardware View** within the Oracle PCA Dashboard provides a graphical representation of the hardware components as they are installed within the rack. The view of the status of these components is static at the time that the page is loaded. A Refresh button at the top of the page allows you to refresh the view at any time.

During particular maintenance tasks, such as upgrading management nodes, you may need to disable compute node provisioning temporarily. This Disable CN Provisioning button at the top of the page allows you to suspend provisioning activity. When compute node provisioning is suspended, the button text

changes to Enable CN Provisioning and its purpose changes to allow you to resume compute node provisioning as required.

Rolling over each item in the graphic with the mouse raises a pop-up window providing the name of the component, its type, and a summary of configuration and status information. For compute nodes, the pop-up window includes a Reprovision button, which allows you to restart the provisioning process if the node becomes stuck in an intermittent state or goes into error status before it is added to the Oracle VM server pool. Instructions to reprovision a compute node are provided in Section 7.12, “A Compute Node Fails to Complete Provisioning”.

Caution

The Reprovision button is to be used *only* for compute nodes that fail to complete provisioning. For compute nodes that have been provisioned properly and/or host running virtual machines, the Reprovision button is made unavailable to prevent incorrect use, thus protecting healthy compute nodes from loss of functionality, data corruption, or being locked out of the environment permanently.

Alongside each installed component within the appliance rack, a status icon provides an indication of the *provisioning status* of the component. The Hardware View does not provide real-time health and status information about active components. Its monitoring functionality is restricted to the provisioning process. When a component has been provisioned completely and correctly, the Hardware View continues to indicate correct operation even if the component should fail or be powered off. See Table 2.1 for an overview of the different status icons and their meaning.

Table 2.1 Table of Hardware Provisioning Status Icons

Icon	Status	Description
	OK	The component is running correctly and has passed all health check operations. Provisioning is complete.
	Warning	The component is running, but has not passed all health check operations and investigation may be required.
	Error	The component is not running and has failed health check operations. Component troubleshooting is required and the component may need to be replaced.

Note

For real-time health and status information of your active Oracle PCA hardware, after provisioning, consult the Oracle VM Manager or Oracle Enterprise Manager UI.

If either of the Fabric Interconnects is clicked on in the Hardware View, the view is changed to the I/O Ports tab in the Network View. See Section 2.3, “Network View” for more information.

The Hardware View provides an accessible tool for troubleshooting hardware components within the Oracle PCA and identifying where these components are actually located within the rack. Where components might need replacing, the new component must take the position of the old component within the rack to maintain configuration.

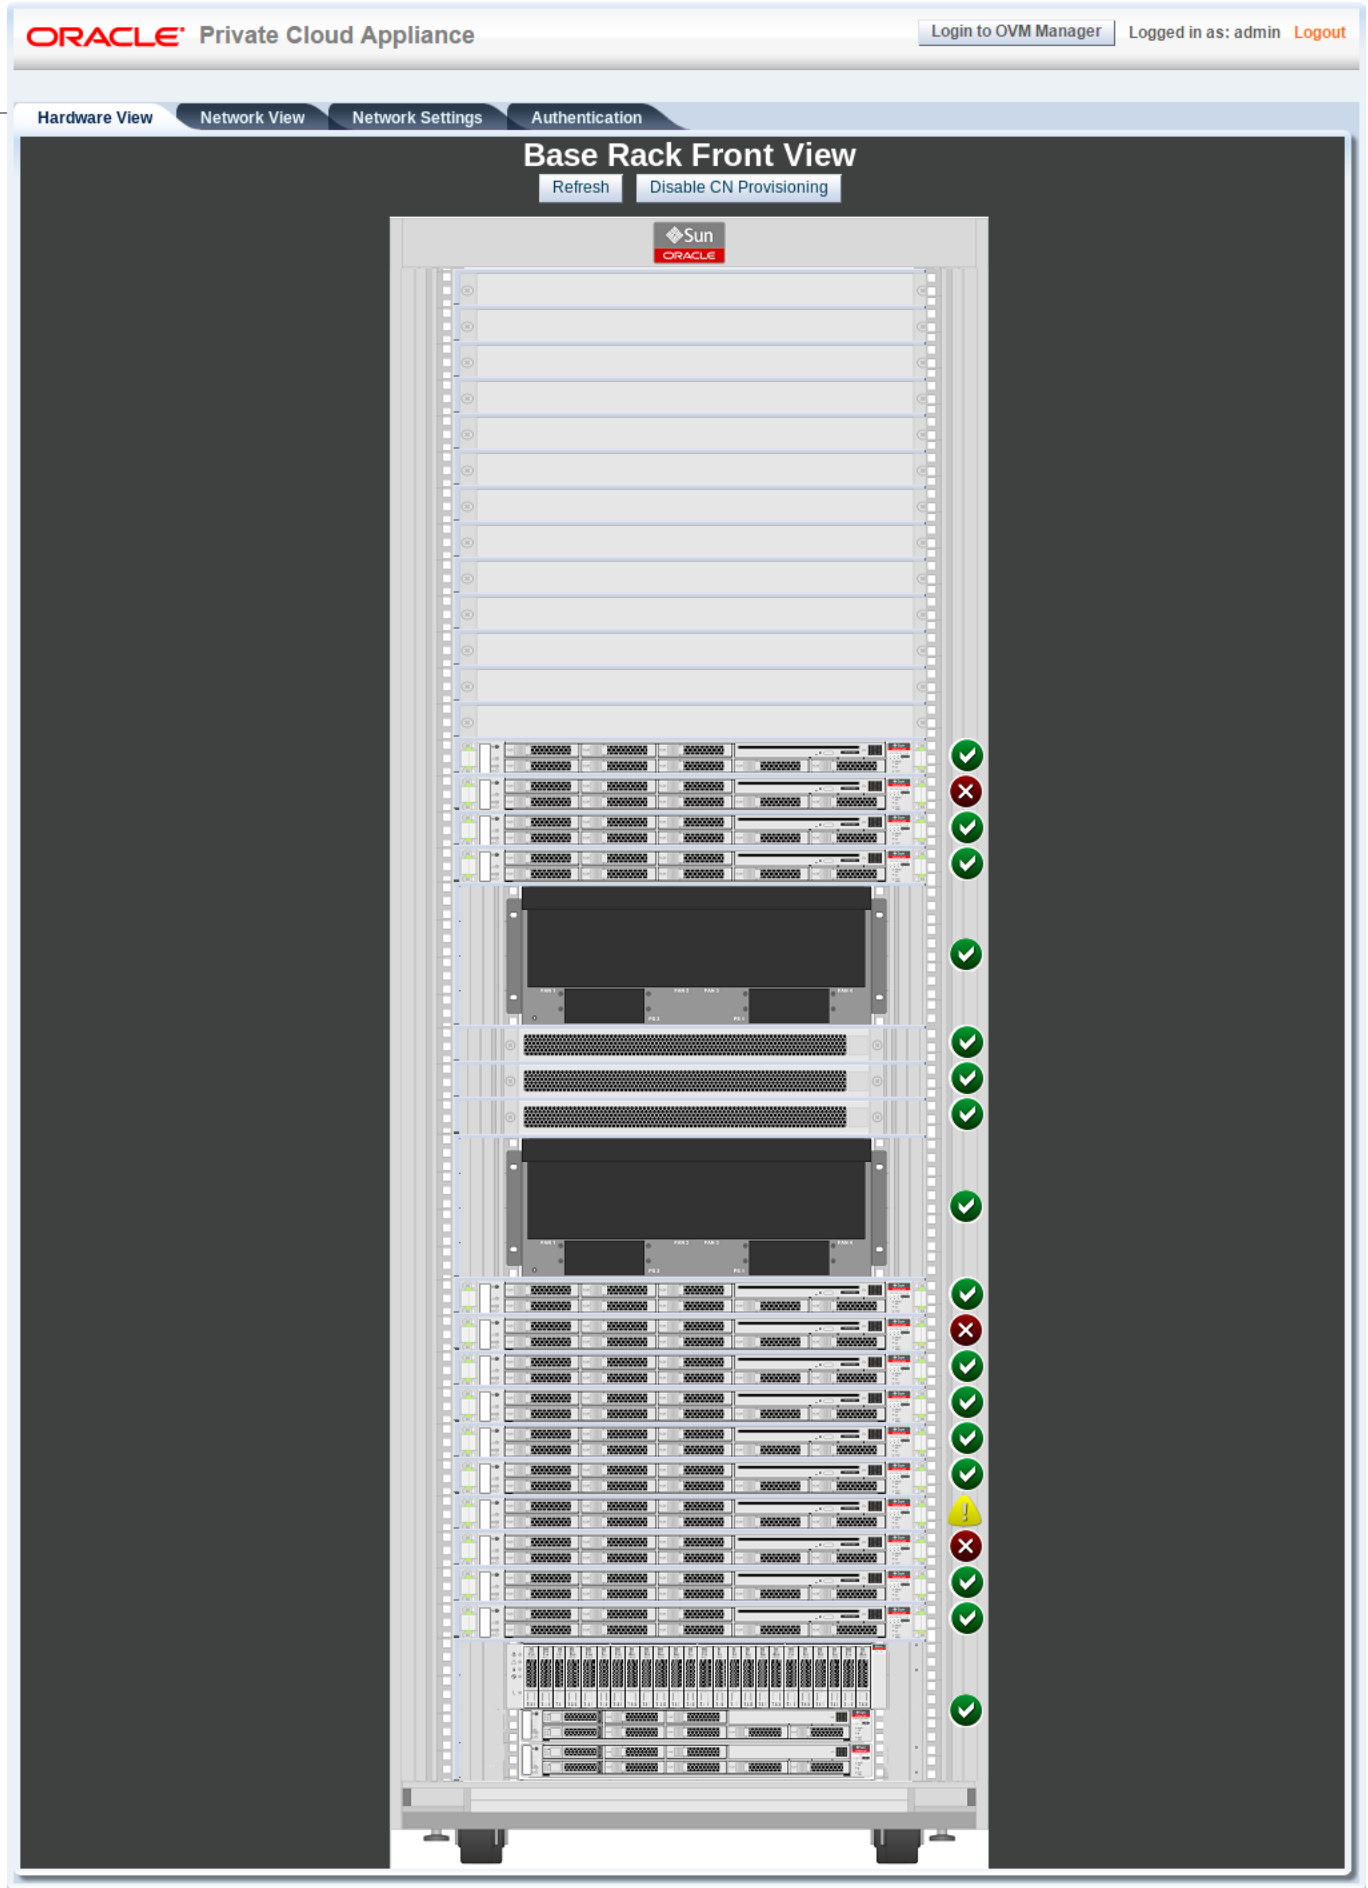
ORACLE Private Cloud Appliance
Login to OVM Manager
Logged in as: admin
Logout
Hardware View
Network View
Network Settings
Authentication
Base Rack Front View
Refresh
Disable CN Provisioning
Sun
ORACLE

2.3 Network View

The **Network View** provides a graphical representation of the networking hardware within the Oracle PCA rack. Specifically, the I/O modules and their network ports on the are shown on a diagram depicting the back panel of the Fabric Interconnect.

Figure 2.2 Network View - IO Module Ports

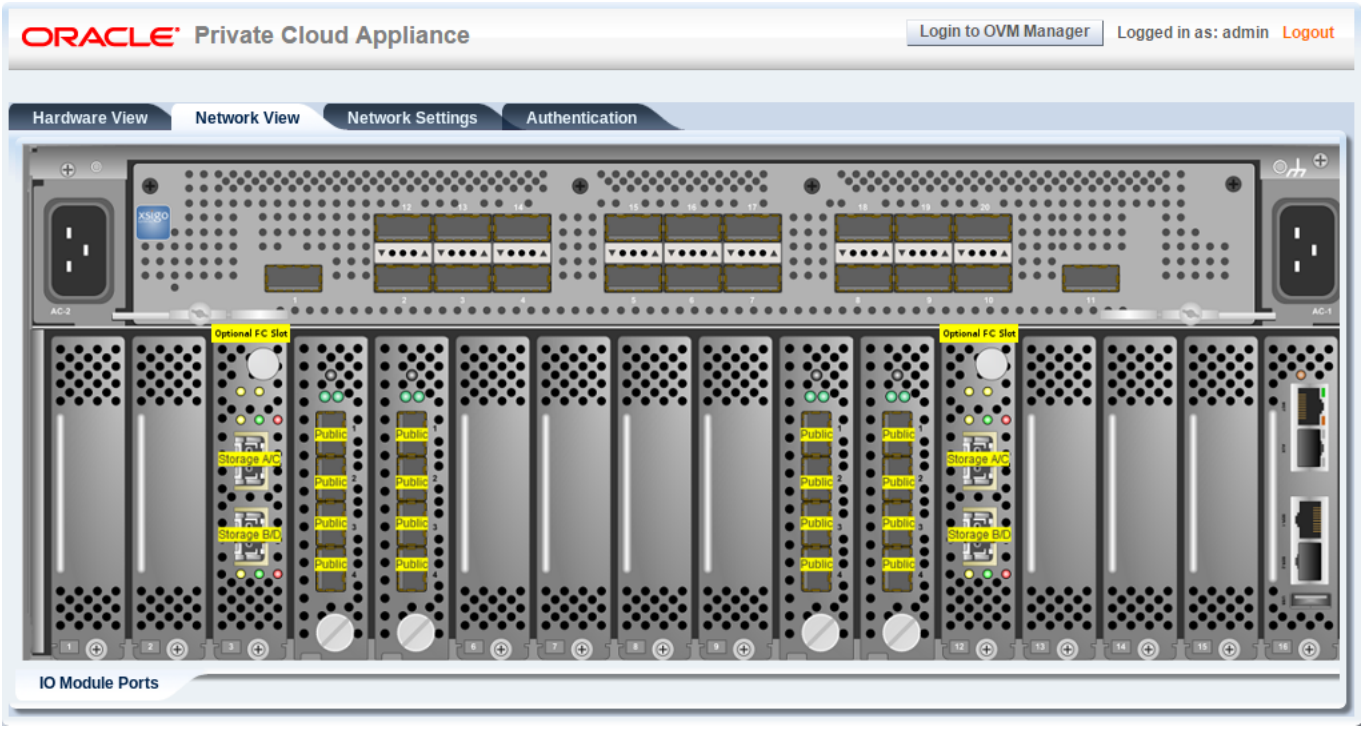

Note that the Fibre Channel (FC) ports on the PCIe Slots 3 and 12 are reserved for use to connect external FC storage to the Oracle PCA . These ports are labeled to indicate the storage clouds that are related to each FC port and can be cabled to one or more external FC switches in the manner described in Adding External Fibre Channel Storage in the *Oracle Private Cloud Appliance Installation Guide*.

2.4 Network Settings

The **Network Settings** tab is used to configure networking and service information for the management nodes. The following fields are available for configuration:

- **Management Node 1:**
  - **IP Address:** Specify an IP address within your datacenter network that can be used to directly access this management node.
  - **Hostname:** Specify the host name for the first management node system.
- **Management Node 2:**
  - **IP Address:** Specify an IP address within your datacenter network that can be used to directly access this management node.
  - **Hostname:** Specify the host name for the second management node system.

- **Management Node Cluster:**
  - **Management Network VLAN:**
    - **No VLAN:** By default this radio button is set. The default configuration does not assume that your management network exists on a VLAN.
    - **VLAN Enabled:** If you have configured a VLAN on your switch for the management network, you should set this radio button and then specify the VLAN ID for the VLAN that the management network must use in the provided field.
  - **Virtual IP:** Specify the shared Virtual IP address that is used to always access the active management node. This IP address must be in the same subnet as the IP addresses that you have specified for each management node.
  - **Domain Name:** Specify the domain that each of the management node systems belong to.
  - **Netmask:** Specify the netmask for the network that the Virtual IP address and management node IP addresses belong to.
  - **Default Gateway:** Specify the default gateway for the network that the Virtual IP address and management node IP addresses belong to.
  - **NTP:** Specify the NTP server that the management nodes and other appliance components must use to synchronize their clocks to.
  - **DNS:** Specify at least one DNS server that the management nodes can use for domain name resolution.
- **Current PCA Password:** You must enter the current Oracle PCA password to make changes to any of these settings.

Clicking the OK button at the bottom of the page saves the settings that are currently displayed on the page and updates the configuration on each of the management nodes. The `ovca` services are restarted in the process, so you are required to log back in to the Dashboard afterwards.

Figure 2.3 Network Settings

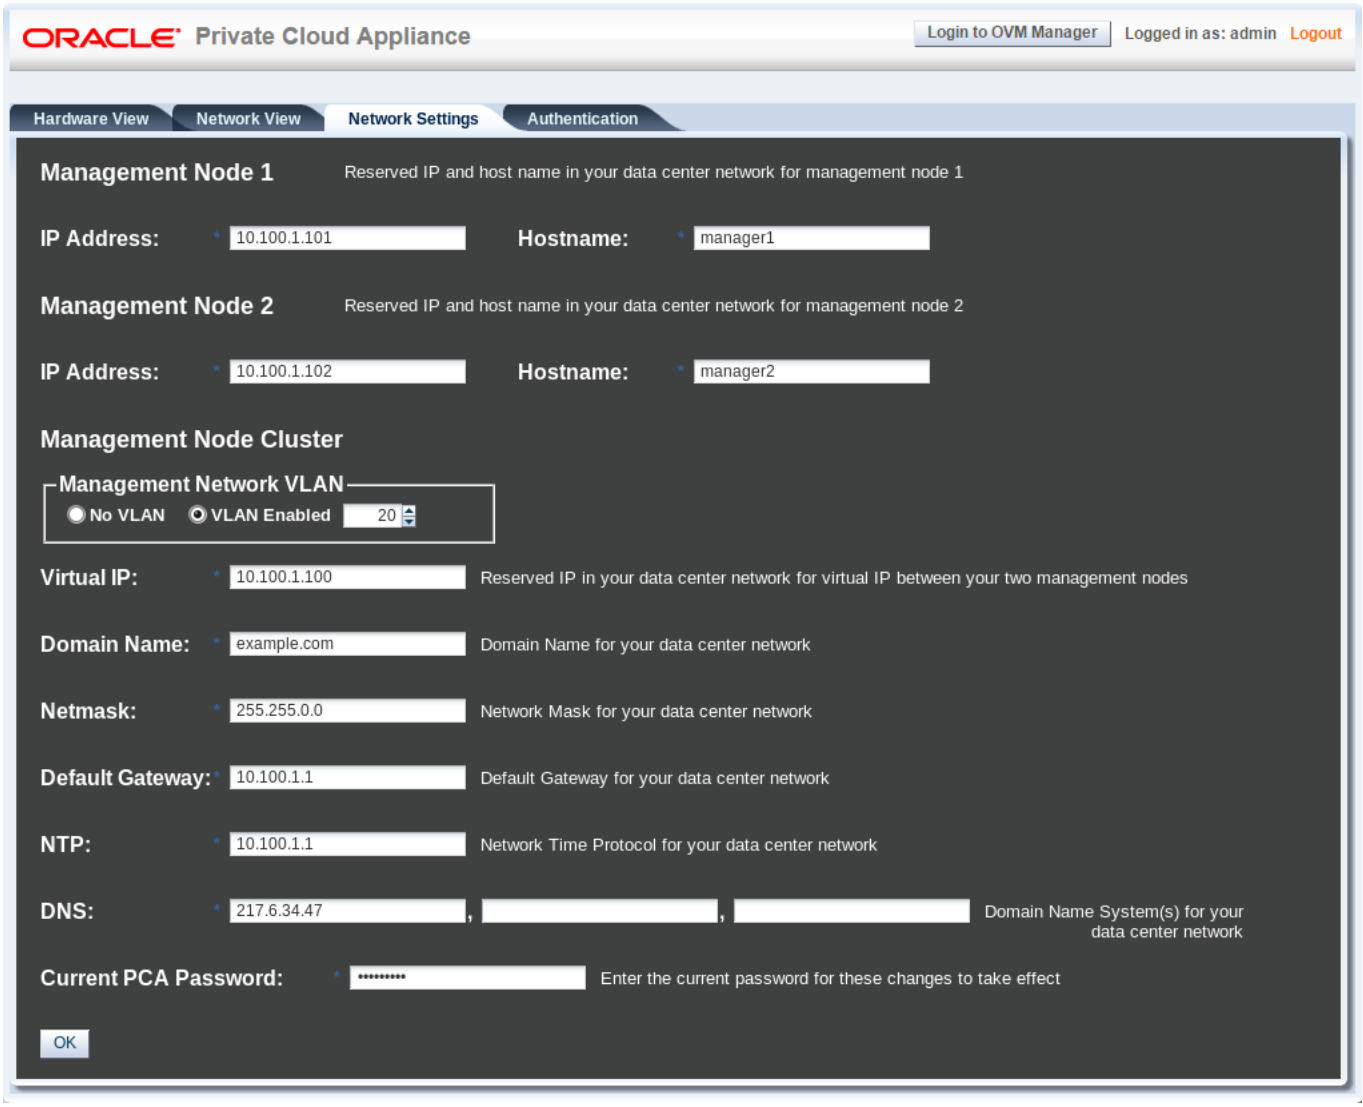

2.5 Authentication

The **Authentication** tab is used to reset the global Oracle PCA password:

- **Current PCA Password:** If changing the Oracle PCA password, it is necessary to provide the current password in this field before specifying the new password in the New PCA Password field.
- **New PCA Password:** Provide a new global Oracle PCA password, to reset the password used for all system configuration components.
- **New PCA Password Confirmation:** Confirm the new Oracle PCA password and check that you have not mis-typed what you intended.

If you used this form to reset the Oracle PCA password, the main appliance configuration components are updated with the same administrative password, except for the Oracle WebLogic Server, which must be changed manually. For information about password management, see Section 1.3.2, “Password Manager (Wallet)”. For instructions to change the Oracle WebLogic Server password, see Section 7.9, “Changing Oracle WebLogic Server Passwords”.

Figure 2.4 Authentication

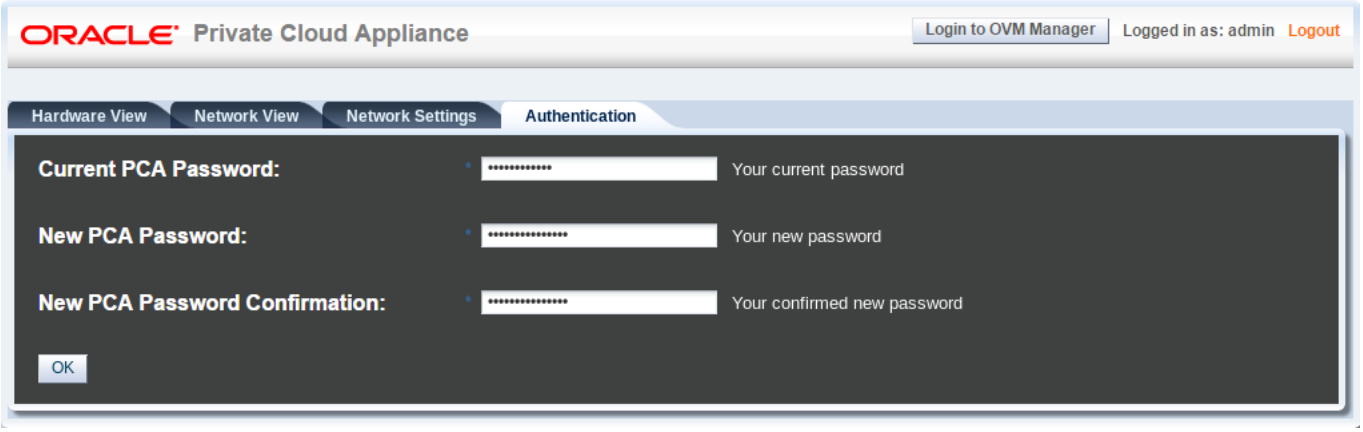

The password update functionality in the Dashboard is restricted to the Dashboard UI itself, Oracle VM Manager, the root user account on both management nodes, and the ovs user account of the Oracle VM mySQL database. It is possible to change the password for other accounts and components by setting the password manually and updating the corresponding Wallet entry.

Caution

ILOM passwords are not modified through the Authentication tab. When changing component ILOM passwords, make sure that it is identical for all compute nodes and management nodes. The server ILOM password must be updated in three locations:

- the ILOM's own interface
- the Oracle PCA Wallet
- the server IPMI configuration in Oracle VM Manager

2.6 Update

Due to the nature of the Oracle Private Cloud Appliance – where the term *appliance* is key – an update is a delicate and complicated procedure that deals with different hardware and software components at the same time. It is virtually impossible to automate the entire process, and more importantly it would be undesirable to take the appliance and the virtual environment it hosts out of service entirely for updating. Instead, updates can be executed in phases and scheduled for minimal downtime. The following table explains how an Oracle PCA update handles different levels or areas of appliance functionality.

Table 2.2 Functional Break-Down of an Appliance Update

Functionality	Physical Location	Description
controller software	management nodes	all components required to set up the management cluster, manage and configure the appliance, and orchestrate compute node provisioning
virtualization platform	compute nodes	all components required to configure the compute nodes and allow virtual machines to be hosted on them
component firmware	infrastructure components	all low-level software components required by the various hardware components for their normal operation as part of the appliance

All the software included in a given release of the Oracle PCA software is tested to work together and should be treated as one package, so no appliance component should be updated individually. All Oracle PCA software releases are downloaded as a single large `.iso` file, which includes the items listed above. Some releases include major updates for one or more appliance components. These are critical stepping stones in a successful upgrade path so they must not be skipped. This is explained in Section 2.6.1, “Mandatory Updates”.

Warning

An update of the Oracle PCA software stack may involve a complete re-imaging of the management nodes. Any customer-installed agents or customizations are overwritten in the process. Before applying new appliance software, back up all local customizations and prepare to re-apply them after the update has completed successfully.

2.6.1 Mandatory Updates

It is expected that not every Oracle PCA system is updated with each minor release. Updates are often selected because they fulfill certain requirements of the environment where the appliance is deployed. However, certain updates bring such significant changes to the product that they must be applied if they are on the path between a system's current software and the intended newer version. This section describes the mandatory updates and the impact they have on the Oracle PCA going forward. A more detailed description of each release can be found in the Oracle Private Cloud Appliance Release Notes.

Release 1.1.3

If your Oracle PCA is currently running Controller Software Release 1.0.2, please upgrade to Release 1.1.3 first. Critical changes include:

- Oracle VM performance tuning
- storage connectivity improvements
- support for a new generation of hardware: Sun Server X4-2 and Oracle ZFS Storage Appliance ZS3-ES
- OpenSSL upgrade to eliminate the 'heartbleed bug' (CVE-2014-0160)

Warning

If you are using a VLAN for your management network and you are upgrading from a previous release where you manually configured your environment, these network changes can adversely affect the upgrade process. Any manual configuration to allow the management network to run on a VLAN must be reverted on both management nodes, before the software process is started.

To update the appliance software to Release 1.1.3 you must first execute the procedure in Section 2.6.2, “Updating the Oracle PCA Controller Software”. After successful completion, perform the additional steps described in these Troubleshooting sections:

- Section 7.5, “Verifying and Re-applying Oracle VM Manager Tuning after Software Update”
- Section 7.6, “Upgrading Existing Compute Node Configuration from Release 1.0.2”

Release 2.0.1

If your Oracle PCA is currently running Controller Software Release 1.1.3, please upgrade to Release 2.0.1 first. Critical changes include:

- external storage support
- appliance management over a VLAN
- Oracle VM upgrade to Release 3.2.8

Warning

If you are using a VLAN for your management network and you are upgrading from a previous release where you manually configured your environment, these network changes can adversely affect the upgrade process. Any manual configuration to allow the management network to run on a VLAN must be reverted on both management nodes, before the software process is started.

As of Release 2.0.1, the Oracle PCA Dashboard provides an option to automatically configure the management network for a VLAN. After the upgrade is completed, you need to set this configuration option to enable VLAN support for the management network. See Section 2.4, “Network Settings” for more information on setting this configuration option.

To update the appliance software to Release 2.0.1 you must first execute the procedure in Section 2.6.2, “Updating the Oracle PCA Controller Software”. After successful completion, perform the additional procedure described in Section 2.6.3, “Upgrading the Virtualization Platform”.

Release 2.0.5

Release 2.0.5 is also a mandatory intermediate step on the upgrade path to a future release. Critical changes include:

- product name change from Oracle Virtual Compute Appliance (VCA) to Oracle Private Cloud Appliance (PCA)
- Oracle Linux kernel upgrade
- support for a new generation of hardware: Oracle Server X5-2 (expansion compute nodes and base racks)
- Oracle VM upgrade to Release 3.2.9

To update the appliance software to Release 2.0.5 you must first execute the procedure in Section 2.6.2, “Updating the Oracle PCA Controller Software”. After successful completion, perform the additional procedure described in Section 2.6.3, “Upgrading the Virtualization Platform”.

2.6.2 Updating the Oracle PCA Controller Software

Updates of the Oracle PCA controller software are performed from the command line of the master management node. Software updates are a three-phase process. First, a zipped ISO containing the updated software must be downloaded from My Oracle Support and made available on an HTTP or FTP server. From there, the ISO is downloaded to the ZFS storage appliance. When the download is complete and the software is unpacked in the appropriate directories, the update is activated and applied to each affected component.

Caution

If direct public access is not available within your data center and you make use of proxy servers to facilitate HTTP, HTTPS and FTP traffic, it may be necessary to edit the Oracle PCA system properties, using the CLI on each management node, to ensure that the correct proxy settings are specified for a download to succeed

from the Internet. This depends on the network location from where the download is served. See Section 7.1, “Adding Proxy Settings for Oracle Private Cloud Appliance Updates” for more information.

Optional: Rebooting the Management Node Cluster

Although not strictly necessary, it may be useful to reboot both management nodes before starting the appliance software update. This leaves the management node cluster in the cleanest possible state, ensures that no system resources are occupied unnecessarily, and eliminates potential interference from processes that have not completed properly.

1. Using SSH and an account with superuser privileges, log into both management nodes using the IP addresses you configured in the Network Setup tab of the Oracle PCA Dashboard. If you use two separate consoles you can view both side by side.

   Note

   The default `root` password is *Welcome1*.

   In all software versions prior to the product name change (Release 2.0.5) you should use `ovca-*` commands instead of `pca-*`.

2. Run the command `pca-check-master` on both management nodes to verify which node owns the master role.

3. Reboot the management node that is **NOT** currently the master. Enter `init 6` at the prompt.

4. Ping the machine you rebooted. When it comes back online, reconnect using SSH and monitor system activity to determine when the secondary management node takes over the master role. Enter this command at the prompt: `tail -f /var/log/messages`. New system activity notifications will be output to the screen as they are logged.

5. In the other SSH console, which is connected to the current active management node, enter `init 6` to reboot the machine and initiate management node failover.

   The log messages in the other SSH console should now indicate when the secondary management node takes over the master role.

6. Verify that both management nodes have come back online after reboot and that the master role has been transferred to the other manager. Run the command `pca-check-master` on both management nodes.

   If this is the case, proceed with the software update procedure below.

Updating the Oracle Private Cloud Appliance Software

Tip

Updates are controlled with `pca-updater` commands. To display usage options, enter `pca-updater --help` at the shell prompt of the management node.

The command used throughout this procedure is `pca-updater`. However, in all software versions prior to the product name change (Release 2.0.5) you should use `ovca-updater` instead.

1. Log into My Oracle Support and download the required Oracle PCA software update.

   You can find the update by searching for the product name “Oracle Private Cloud Appliance”, or for the Patch or Bug Number associated with the update you need.

2. Make the update, a zipped ISO, available on an HTTP or FTP server that is reachable from your Oracle PCA.

3. Using SSH and an account with superuser privileges, log into the management node.

Note

The default `root` password is *Welcome1*.

4. Connect to the management node using its IP address in the data center network, as you configured it in the Network Setup tab of the Oracle PCA Dashboard. For details, see Section 2.4, “Network Settings”.

Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~]#
```

5. Download the ISO to your Oracle PCA.

```
# pca-updater -m download -s -u http://download-url.example.com/pca-version-build.iso.zip
Download of http://download-url.example.com/pca-version-build.iso.zip started successfully. \
Task id is 5ba4f1a09c9546b29a560e3d6682ce67
```

6. Check the progress of the ISO download. When the download is finished, proceed with the next step.

```
# pca-updater -m download -l -i 5ba4f1a09c9546b29a560e3d6682ce67
Task ID                          Pct   Elapsed time Status
5ba4f1a09c9546b29a560e3d6682ce67  60%   155.305 sec    RUNNING
# pca-updater -m download -l -i 5ba4f1a09c9546b29a560e3d6682ce67
Task ID                          Pct   Elapsed time Status
5ba4f1a09c9546b29a560e3d6682ce67  70%   285.904 sec    RUNNING
# pca-updater -m download -l -i 5ba4f1a09c9546b29a560e3d6682ce67
Task ID                          Pct   Elapsed time Status
5ba4f1a09c9546b29a560e3d6682ce67 100%   338.102 sec    SUCCESS
```

To list all downloads and their status, enter this command:

```
# pca-updater -m download -l
Task ID                          Pct   Elapsed time Status
5ba4f1a09c9546b29a560e3d6682ce67 100%   338.102 sec    SUCCESS
```

7. When the download has completed successfully, activate it by launching the update process.

```
# pca-updater -m update -s
WARNING:root:Activating update process. Starting with nonmaster node 192.168.4.4
```

Once you issue this command, the update process is initiated as described in Section 1.7, “Oracle Private Cloud Appliance Software Update”.

8. Check the progress of the software update.

```
# pca-updater -m update -l
Mgmt Node IP    Update Started      Update Ended        Elapsed
192.168.4.4     2013-10-02 16:52:14 ------------------- 0:02:18
```

If you need to cancel and restart the update process, enter this command on the active management node: `pca-updater -m update -x -i 192.168.4.4`.

Note

At a certain point during the update process, the active management node is rebooted. As a result, the SSH connection is lost. In addition, this may cause the Dashboard to become unresponsive temporarily, and you may be required to log back in.

When the master management node reboots, the secondary (updated) management node assumes the master role. The original master management node is then also updated and becomes the backup management node.

Caution

When the first of the two management nodes is updated, the server is configured with a temporary password for approximately one hour: *XXXPASSWORDXXX*. Attempts to log in with the regular password stored in the Wallet will fail during that time. When the update process is complete, the password is reset to the correct value.

2.6.3 Upgrading the Virtualization Platform

Some releases of the Oracle PCA Controller Software include a new version of Oracle VM. As part of the controller software update, the new Oracle VM Manager Release is automatically installed on both management nodes, and a new YUM repository configuration is pushed to all compute nodes to point them to the packages of the new Oracle VM Server Release. As a result, all compute nodes in the Oracle PCA rack are marked as "out of date" in Oracle VM Manager. The *Server Out Of Date* event is informational and does not prevent any operations, but it is recommended that you upgrade all compute nodes to the new Oracle VM Server Release at your earliest convenience.

The Oracle VM Server upgrade was intentionally decoupled from the automated controller software update process. This allows you to plan the compute node upgrades and the migration or downtime of your virtual machines in steps and outside peak hours. As a result, service interruptions for users of the Oracle VM environment can be minimized or even eliminated. By following the instructions in this section, you also make sure that previously deployed virtual machines remain fully functional when the appliance update to the new software release is complete.

During an upgrade of Oracle VM Server, no virtual machine can be running on a given compute node. VMs using resources on a shared storage repository can be migrated to other running compute nodes. If a VM uses resources local to the compute node you want to upgrade, it must be shut down, and returned to service after the Oracle VM Server upgrade.

Upgrading a Compute Node to a Newer Oracle VM Server Release

Caution

Execute this procedure on each compute node *after* the software update on the management nodes has completed successfully. At that stage, the YUM repository configuration of all compute nodes has been updated to point to the new Oracle VM Server packages on the Oracle PCA internal storage.

Note

A large part of this procedure must be executed from Oracle VM Manager. An overview of its usage in the context of Oracle PCA, and links to detailed instructions in the Oracle VM documentation, are provided in this Administrator's Guide, in Chapter 4, *Managing the Oracle VM Virtual Infrastructure*.

1. Make sure that the appliance software has been updated successfully to the new release.

   You can verify this by logging into the master management node and entering the following command:

   ```
   # pca-diag version
   2.0.5_54_2015-06-11
   ```

2. Log in to Oracle VM Manager.

   For details, see Section 4.2, “Logging in to the Oracle VM Manager Web UI”.

3. Identify the master server in the server pool.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Info** perspective from the list in the toolbar.

      A configuration overview of the server pool is displayed.

   c. Locate the *Master Server* setting and write down which server is the pool master.

   Caution

   The server pool master server should be the last Oracle VM Server to go through this upgrade procedure. This allows you to upgrade the compute nodes one by one without causing the master role to be transferred to another node in the server pool each time. This eliminates potential resource locks and errors related to transferring the master role, and saves a significant amount of time.

4. Migrate all running virtual machines away from the compute node you want to upgrade.

   Information on migrating virtual machines is provided in the Oracle VM User's Guide section entitled Migrating Virtual Machines.

5. Place the compute node in maintenance mode.

   Information on maintenance mode is provided in the Oracle VM User's Guide section entitled Placing an Oracle VM Server into Maintenance Mode.

   a. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

      The **Edit Server** dialog box is displayed.

   b. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

      The Oracle VM Server is in maintenance mode and ready for servicing.

6. Run the Oracle VM Server update on the compute node in question.

   Information on the Oracle VM Server update functionality is provided in the Oracle VM User's Guide section entitled Updating and Upgrading Oracle VM Servers.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Servers** perspective from the list in the toolbar.

A table is displayed with all servers in the server pool.

c. Select the appropriate compute node in the management pane table and click **Update Server** in the management pane toolbar.

As part of the update procedure, the Oracle VM Server is restarted but remains in maintenance mode.

Warning

If the compute node does not reboot during the update, you must restart it from within Oracle VM Manager.

7. Take the compute node out of maintenance mode.

a. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

b. Clear the **Maintenance Mode** check box. Click OK.

The Oracle VM Server rejoins the server pool as a fully functioning member.

8. Repeat this procedure for each compute node in your Oracle PCA.

Caution

The server pool master server should be the last Oracle VM Server to go through this upgrade procedure. This allows you to upgrade the compute nodes one by one without causing the master role to be transferred to another node in the server pool each time. This eliminates potential resource locks and errors related to transferring the master role, and saves a significant amount of time.

The appliance software update is now complete. The Oracle PCA is ready to resume all normal operations.

2.6.4 Upgrading Component Firmware

All the software components in a given Oracle PCA release are designed to work together. As a general rule, no individual appliance component should be upgraded. If a firmware upgrade is required for one or more components, the correct version is distributed inside the Oracle PCA `.iso` file you downloaded from My Oracle Support. When the image file is unpacked on the internal shared storage, the firmwares are located in this directory: `/nfs/shared_storage/mgmt_image/firmware/` .

There are currently no specific or additional procedures to upgrade the firmware of an Oracle PCA hardware component. Unless otherwise instructed, you may follow the instructions provided in the product documentation of the subcomponent. An overview of the documentation for appliance components can be found in the Preface of this book and on the index page of the Oracle PCA Documentation Library.

Chapter 3 The Oracle Private Cloud Appliance Command Line Interface (CLI)

Table of Contents

3.1 CLI Usage 37
- 3.1.1 Interactive Mode 38
- 3.1.2 Single-command Mode 39
- 3.1.3 Controlling CLI Output 40
- 3.1.4 Internal CLI Help 42

3.2 CLI Commands 43
- 3.2.1 backup 43
- 3.2.2 diagnose 44
- 3.2.3 delete 45
- 3.2.4 list 46
- 3.2.5 reprovision 50
- 3.2.6 rerun 50
- 3.2.7 set 51
- 3.2.8 show 52
- 3.2.9 start 56
- 3.2.10 stop 56

All Oracle Private Cloud Appliance (PCA) command line utilities are consolidated into a single command line interface that is accessible via the management node shell by running the `pca-admin` command located at `/usr/sbin/pca-admin`. This command is in the system path for the root user, so you should be able to run the command from anywhere that you are located on a management node. The CLI provides access to all of the tools available in the Oracle PCA Dashboard, as well as many that do not have a Dashboard equivalent. The design of the CLI makes it possible to script actions that may need to be performed more regularly, or to write integration scripts with existing monitoring and maintenance software not directly hosted on the appliance.

It is important to understand that the CLI, described here, is distinct from the Oracle VM Manager command line interface, which is described fully in the Oracle VM documentation available at http://docs.oracle.com/cd/E35328_01/E35336/html/index.html.

In general, it is preferable that CLI usage is restricted to the active management node. While it is possible to run the CLI from either management node, some commands are restricted to the active management node and return an error if you attempt to run them on the passive management node.

3.1 CLI Usage

The Oracle PCA command line interface is triggered by running the `pca-admin` command. It can run either in interactive mode (see Section 3.1.1, “Interactive Mode”) or in single-command mode (see Section 3.1.2, “Single-command Mode”) depending on whether you provide the syntax to run a particular CLI command when you invoke the command line interpreter.

The syntax when using the CLI is as follows:

```
PCA> Command Command_Target <Arguments> Options
```

where:

- *Command* is the command type that should be initiated. For example `list`;
- *Command_Target* is the Oracle PCA component or process that should be affected by the command. For example `management-node`, `compute-node`, `task` etc;
- *<Arguments>* consist of positioning arguments related to the command target. For instance, when performing a reprovisioning action against a compute node, you should provide the specific compute node that should be affected as an argument for this command. For example: `reprovision compute-node ovcacn11r1`;
- *Options* consist of options that may be provided as additional parameters to the command to affect its behavior. For instance, the `list` command provides various sorting and filtering options that can be appended to the command syntax to control how output is returned. For example: `list compute-node --filter-column ILOM_State --filter dead`. See Section 3.1.3, “Controlling CLI Output” for more information on many of these options.

The CLI includes its own internal help that can assist you with understanding the commands, command targets, arguments and options available. See Section 3.1.4, “Internal CLI Help” for more information on how to use this help system. When used in interactive mode, the CLI also provides tab completion to assist you with the correct construction of a command. See Section 3.1.1.1, “Tab Completion” for more information on this.

3.1.1 Interactive Mode

The Oracle PCA command line interface (CLI) provides an interactive shell that can be used for user-friendly command line interactions. This shell provides a closed environment where users are able to enter commands specific to the management of the Oracle PCA. By using the CLI in interactive mode, the user can avail of features like tab completion to easily complete commands correctly. By default, running the `pca-admin` command without providing any additional parameters causes the CLI interpreter to run in interactive mode.

It is possible to identify that you are in a CLI shell running in interactive mode as the shell prompt is indicated by **PCA>**.

Example 3.1 An example of interactive mode usage of the CLI

```
# pca-admin
Welcome to PCA! Release: 2.0.5
PCA>
PCA> list management-node

Management_Node IP_Address    Node_State    ILOM_MAC             ILOM_State    Master
--------------- ----------    ----------    --------             ----------    ------
ovcamn06r1      192.168.4.4   RUNNING       00:10:e0:2e:7c:53    running
ovcamn05r1      192.168.4.3   RUNNING       00:10:e0:31:93:6f    running       Yes
----------------
2 rows displayed

Status: Success
PCA> exit
#
```

To exit from the CLI when it is in interactive mode, you can use either the `quit` or `exit` command, or alternatively use the Ctrl+D key combination.

3.1.1.1 Tab Completion

The CLI supports tab-completion when in interactive mode. This means that pressing the tab key while entering a command can either complete the command on your behalf, or can indicate options and

possible values that can be entered to complete a command. Usually you must press the tab key at least twice to effect tab-completion.

Tab-completion is configured to work at all levels within the CLI and is context sensitive. This means that you can press the tab key to complete or prompt for commands, command targets, options, and for certain option values. For instance, pressing the tab key twice at a blank prompt within the CLI automatically lists all possible commands, while pressing the tab key after typing the first letter or few letters of a command automatically completes the command for you. Once a command is specified, pressing the tab key indicates command targets. If you have specified a command target, pressing the tab key indicates other options available for the command sequence. If you press the tab key after specifying a command option that requires an option value, such as the `--filter-column` option, the CLI attempts to provide you with the values that can be used with that option.

Example 3.2 Examples showing tab-completion

```
PCA> <tab>
EOF          backup       delete       diagnose     exit         help
list         q            quit         reprovision  rerun        set
shell        show         start        stop

PCA> list <tab>
compute-node     config-error     management-node  network-cloud
opus-port        storage-cloud    task             wwpn-info

PCA> list m<tab>anagement-node
```

The *<tab>* indicates where the user pressed the tab key while in an interactive CLI session. In the final example, the command target is automatically completed by the CLI.

3.1.1.2 Running Shell Commands

It is possible to run standard shell commands while you are in the CLI interpreter shell. These can be run by either preceding them with the `shell` command or by using the `!` operator as a shortcut to indicate that the command that follows is a standard shell command. For example:

```
PCA> shell date
Wed Jun 24 04:54:54 PDT 2015
PCA> !uptime > /tmp/uptime-today
PCA> !rm /tmp/uptime-today
```

3.1.2 Single-command Mode

The CLI supports 'single-command mode', which allows you to execute a single command from the shell via the CLI and to obtain the output before the CLI exits back to the shell. This is particularly useful when writing scripts that may interact with the CLI, particularly if used in conjunction with the CLI's JSON output mode described in Section 3.1.3.1, “JSON Output”.

To run the CLI in single-command mode, simply include the full command syntax that you wish to execute as parameters to the `pca-admin` command.

An example of single command mode is provided below:

```
# pca-admin list compute-node

Compute_Node    IP_Address      Node_State      ILOM_MAC            ILOM_State
------------    ----------      ----------      --------            ----------
ovcacn08r1      192.168.4.9     RUNNING         00:10:e0:65:2f:b7   running
ovcacn28r1      192.168.4.10    RUNNING         00:10:e0:62:31:81   running
ovcacn10r1      192.168.4.7     RUNNING         00:10:e0:65:2f:cf   running
```

```
ovcacn30r1      192.168.4.8     RUNNING         00:10:e0:40:cb:59    running
ovcacn07r1      192.168.4.11    RUNNING         00:10:e0:62:ca:09    running
ovcacn26r1      192.168.4.12    RUNNING         00:10:e0:65:30:f5    running
ovcacn29r1      192.168.4.5     RUNNING         00:10:e0:31:49:1d    running
ovcacn09r1      192.168.4.6     RUNNING         00:10:e0:65:2f:3f    running
----------------
8 rows displayed
Status: Success

#
```

3.1.3 Controlling CLI Output

The CLI provides options to control how output is returned in responses to the various CLI commands that are available. These are provided as additional options as the final portion of the syntax for a CLI command. Many of these options can make it easier to identify particular items of interest through sorting and filtering, or can be particularly useful when scripting solutions as they help to provide output that is more easily parsed.

3.1.3.1 JSON Output

JSON format is a commonly used format to represent data objects in a way that is easy to machine-parse but is equally easy for a user to read. Although JSON was originally developed as a way to represent JavaScript objects, parsers are available for a wide number of programming languages, making it an ideal output format for the CLI if you are scripting a custom solution that may need to interface directly with the CLI.

The CLI returns its output for any command in JSON format if the `--json` option is specified when a command is run. Typically this option may be used when running the CLI in single-command mode. An example follows:

```
# pca-admin list compute-node --json
{
    "00:10:e0:65:2f:b2": {
        "ip": "192.168.4.9",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn08r1",
        "ilom_mac": "00:10:e0:65:2f:b7"
    },
    "00:10:e0:62:31:7c": {
        "ip": "192.168.4.10",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn28r1",
        "ilom_mac": "00:10:e0:62:31:81"
    },
    "00:10:e0:65:2f:ca": {
        "ip": "192.168.4.7",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn10r1",
        "ilom_mac": "00:10:e0:65:2f:cf"
    },
    "00:10:e0:40:cb:54": {
        "ip": "192.168.4.8",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn30r1",
        "ilom_mac": "00:10:e0:40:cb:59"
    },
    "00:10:e0:62:ca:04": {
        "ip": "192.168.4.11",
```

```
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn07r1",
        "ilom_mac": "00:10:e0:62:ca:09"
    },
    "00:10:e0:65:30:f0": {
        "ip": "192.168.4.12",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn26r1",
        "ilom_mac": "00:10:e0:65:30:f5"
    },
    "00:10:e0:31:49:18": {
        "ip": "192.168.4.5",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn29r1",
        "ilom_mac": "00:10:e0:31:49:1d"
    },
    "00:10:e0:65:2f:3a": {
        "ip": "192.168.4.6",
        "state": "RUNNING",
        "ilom_state": "running",
        "name": "ovcacn09r1",
        "ilom_mac": "00:10:e0:65:2f:3f"
    }
}
```


In some cases the JSON output may contain more information than is displayed in the tabulated output that is usually shown in the CLI when the `--json` option is not used. Furthermore, the keys used in the JSON output may not map identically to the table column names that are presented in the tabulated output.

Sorting and filtering options are currently not supported in conjunction with JSON output, since these facilities can usually be implemented on the side of the parser.

3.1.3.2 Sorting

Typically, when using the `list` command, you may wish to sort information in a way that makes it easier to view items of particular interest. This is achieved using the `--sorted-by` and `--sorted-order` options in conjunction with the command. When using the `--sorted-by` option, you must specify the column name against which the sort should be applied. You can use the `--sorted-order` option to control the direction of the sort. This option should be followed either with ASC for an ascending sort, or DESC for a descending sort. If this option is not specified, the default sort order is ascending.

For example, to sort a view of compute nodes based on the status of the ILOM for each compute node, you may do the following:

```
PCA> list compute-node --sorted-by ILOM_State --sorted-order ASC

Compute_Node    IP_Address      Node_State      ILOM_MAC            ILOM_State
------------    ----------      ----------      --------            ----------
ovcacn08r1      192.168.4.9     RUNNING         00:10:e0:65:2f:b7   dead
ovcacn28r1      192.168.4.10    RUNNING         00:10:e0:62:31:81   initializing_stage_wait_for_hmp
ovcacn10r1      192.168.4.7     RUNNING         00:10:e0:65:2f:cf   initializing_stage_wait_for_hmp
ovcacn30r1      192.168.4.8     RUNNING         00:10:e0:40:cb:59   running
ovcacn07r1      192.168.4.11    RUNNING         00:10:e0:62:ca:09   running
ovcacn26r1      192.168.4.12    RUNNING         00:10:e0:65:30:f5   running
ovcacn29r1      192.168.4.5     RUNNING         00:10:e0:31:49:1d   running
ovcacn09r1      192.168.4.6     RUNNING         00:10:e0:65:2f:3f   running
----------------
8 rows displayed

Status: Success
```

Note that the `--sorted-by` option is case sensitive, but you can use tab-completion to easily obtain the options for different column names. See Section 3.1.1.1, “Tab Completion” for more information.

3.1.3.3 Filtering

Some tables may contain a large number of rows that you are not interested in, to limit the output to items of particular interest you can use the filtering capabilities that are built into the CLI. Filtering is achieved using a combination of the `--filter-column` and `--filter` options. The `--filter-column` option must be followed by specifying the column name, while the `--filter` option is followed with the specific text that should be matched to form the filter. The text that should be specified for a `--filter` must be an exact match. Filtering does not currently support case insensitivity, wildcards, regular expressions or partial matches.

For example, to view only the compute nodes that have an ILOM status equivalent to 'dead', you could use the following filter:

```
PCA> list compute-node --filter-column ILOM_State --filter dead

Compute_Node    IP_Address      Node_State      ILOM_MAC            ILOM_State
------------    ----------      ----------      --------            ----------
ovcacn09r1      192.168.4.10    DEAD            00:10:e0:0f:55:cb   dead
ovcacn11r1      192.168.4.9     DEAD            00:10:e0:0f:57:93   dead
ovcacn14r1      192.168.4.7     DEAD            00:10:e0:46:9e:45   dead
ovcacn36r1      192.168.4.11    DEAD            00:10:e0:0f:5a:9f   dead
----------------
4 rows displayed

Status: Success
```

Note that the `--filter-column` option is case sensitive, but you can use tab-completion to easily obtain the options for different column names. See Section 3.1.1.1, “Tab Completion” for more information.

3.1.4 Internal CLI Help

The CLI includes its own internal help system. This is triggered by issuing the `help` command:

```
PCA> help

Documented commands (type help <topic>):
========================================
backup  diagnose  list         rerun  show   stop
delete  help      reprovision  set    start

Undocumented commands:
======================
EOF  exit  q  quit  shell
```

The help system displays all of the available commands that are supported by the CLI. These are organized into 'Documented commands' and 'Undocumented commands'. Undocumented commands are usually commands that are not specific to the management of the Oracle PCA, but are mostly discussed within this documentation. Note that more detailed help can be obtained for any documented command by appending the name of the command to the `help` query. For example, to obtain the help documentation specific to the list command, you can do the following:

```
PCA> help list

Usage: pca-admin list <Command Target> [OPTS]

Command Targets:
  storage-cloud:        List storage clouds.
  task:                 List task.
```

```
  management-node:      List management node.
  opus-port:            List OPUS port.
  config-error:         List configuration errors.
  compute-node:         List computer node.
  network-cloud:        List network clouds.
  wwpn-info:            List wwpn info.

Options:
  --json                Display the output in json format.
  --sorted-by=SORTEDBY  Sorting the table by a column.
  --sorted-order=SORTEDORDER
                        Sorting order.
  --filter-column=FILTERCOLUMN
                        Table column that needs to be filtered.
  --filter=FILTER       filter criterion
```

You can drill down further into the help system for most commands by also appending the command target onto your `help` query:

```
PCA> help reprovision compute-node

  Usage:
  reprovision compute-node <compute node name> [options]

  Example:
  reprovision compute-node ovcacn11r1

  Description:
  Reprovision a compute node.
```

Finally, if you submit a help query for something that doesn't exist, the help system generates an error and automatically attempts to prompt you with alternative candidates:

```
PCA> list ta
Status: Failure
Error Message: Error (MISSING_TARGET_000): Missing command target for command: list.
Command targets can be: ['storage-cloud', 'task', 'management-node', 'opus-port',
'config-error', 'compute-node', 'network-cloud', 'wwpn-info'].
```

3.2 CLI Commands

This section describes all of the documented commands available via the CLI.

3.2.1 backup

Triggers a manual backup of the Oracle Private Cloud Appliance.

Syntax

`backup` [`--json`]

Description

Use the `backup` command to initiate a backup task outside of the usual cron schedule. The backup task performs a full backup of the Oracle PCA as described in Section 1.6, “Oracle Private Cloud Appliance Backup”. The CLI command does not monitor the progress of the backup task itself, and exits immediately after triggering the task, returning the task ID, its initial status, its progress and start time. This command must only ever be run on the active management node.

You can use the `show task` command to view the status of the task after you have initiated the backup. See Example 3.21, “Show Task” for more information.

Options

The following table shows the available options for this command.

Option	Description
`--json`	Return the output of the command in JSON format

Examples

Example 3.3 Running a backup task

```
PCA> backup

Task_ID                          Status  Progress Start_Time
-------                          ------  -------- ----------
9605d0980c974ee38ef7de8eb6220d74  RUNNING        0 06-02-2014 07:20:32
---------------
1 row displayed

Status: Success
```

3.2.2 diagnose

Performs various diagnostic checks against the Oracle Private Cloud Appliance for support purposes.

Syntax

`diagnose` { `ilom` | `software` } [`--json`]

Description

Use the `diagnose` command to initiate a diagnostic check of various components that make up Oracle PCA.

Options

The following table shows the available options for this command.

Option	Description
{ `ilom` | `software` }	Specify the type of diagnostic check that should be performed as the command target. The `ilom` diagnostic checks that the ILOM for each component is accessible on the management network. The `software` diagnostic triggers the Oracle PCA software acceptance tests.
`--json`	Return the output of the command in JSON format

Examples

Example 3.4 Running the ILOM Diagnostic

```
PCA> diagnose ilom
Checking ILOM health............please wait..

IP_Address      Status          Health_Details
----------      ------          --------------
192.168.4.129   Not Connected
```

```
192.168.4.128   Not Connected
192.168.4.127   Not Connected
192.168.4.126   Not Connected
192.168.4.125   Not Connected
192.168.4.124   Not Connected
192.168.4.123   Not Connected
192.168.4.122   Not Connected
192.168.4.121   Not Connected
192.168.4.120   Not Connected
192.168.4.101   OK
192.168.4.102   OK
192.168.4.105   Faulty          Mon Nov 25 14:17:37 2013  Power    PS1 (Power Supply 1)
                                A loss of AC input to a power supply has occurred.
                                (Probability: 100, UUID: 2c1ec5fc-ffa3-c768-e602-ca12b86e3ea1,
                                Part Number: 07047410, Serial Number: 476856F+1252CE027X,
                                Reference Document: http://www.sun.com/msg/SPX86-8003-73)
192.168.4.107   OK
192.168.4.106   OK
192.168.4.109   OK
192.168.4.108   OK
192.168.4.112   OK
192.168.4.113   Not Connected
192.168.4.110   OK
192.168.4.111   OK
192.168.4.116   Not Connected
192.168.4.117   Not Connected
192.168.4.114   Not Connected
192.168.4.115   Not Connected
192.168.4.118   Not Connected
192.168.4.119   Not Connected
-----------------
27 rows displayed

Status: Success
```

Example 3.5 Running the Software Diagnostic

```
PCA> diagnose software
PCA Software Acceptance Test runner utility
Test -    701 - OpenSSL CVE-2014-0160 Heartbleed bug Acceptance [PASSED]
Test -    785 - PCA package Acceptance [PASSED]
Test -   1083 - Mgmt node xsigo network interface Acceptance [PASSED]
Test -    787 - Shared Storage Acceptance [PASSED]
Test -    973 - Simple connectivity Acceptance [PASSED]
Test -   1078 - Test for ovs-agent service on CNs Acceptance [PASSED]
Test -   1079 - Test for shares mounted on CNs Acceptance [PASSED]
Test -   1080 - ovs-log check Acceptance [PASSED]
Test -    788 - PCA services Acceptance [PASSED]
Test -    789 - PCA config file Acceptance [PASSED]
Test -   1300 - All compute nodes running Acceptance [PASSED]
Test -   1318 - Check support packages in PCA image Acceptance [PASSED]
Test -    928 - Repositories defined in OVM manager Acceptance [PASSED]
Test -   1107 - Compute node xsigo network interface Acceptance [PASSED]
Test -   1316 - PCA version Acceptance [PASSED]
Test -   1117 - Network interfaces check Acceptance [PASSED]
Test -    824 - OVM manager settings Acceptance [PASSED]
Test -    927 - OVM server model Acceptance [PASSED]
Test -    925 - PCA log Acceptance [PASSED]
Test -    926 - Networks defined in OVM manager for CNs Acceptance [PASSED]
Test -    822 - Compute node network interface Acceptance [PASSED]
Status: Success
```

3.2.3 delete

The `delete` command can be used to delete a failed configuration task from the configuration error database.

Syntax

```
delete { config-error } id [ --confirm ] [ --json ]
```

where `id` is the identifier for the configuration error for the item that you wish to delete from the database.

Description

Use the `delete` command to remove a configuration error or a task from the respective database. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Once a configuration error has been deleted from the database, you may not be able to re-run the configuration task associated with it. To obtain a list of configuration errors, use the `list config-error` command. See Example 3.15, “List All Configuration Errors” for more information.

Options

The following table shows the available options for this command.

Option	Description
{ `config-error` | `task` }	The command target to perform the delete operation against.
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--json`	Return the output of the command in JSON format

Examples

Example 3.6 Removing a Configuration Error

```
PCA> delete config-error 87
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y
Status: Success
```

Example 3.7 Removing a Task

```
PCA> delete task 1182919afe584c508fd87b02211d6760
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y
Status: Success
```

3.2.4 list

The `list` command can be used to list the different components and tasks within the Oracle PCA. The output displays information relevant to each component or task. Output from the list command is usually tabulated so that different fields appear as columns for each row of information relating to the command target.

Syntax

`list` { `compute-node` | `config-error` | `management-node` | `network-cloud` | `opus-port` | `storage-cloud` | `task` | `wwpn-info` } [--json] [[--sorted-by `SORTEDBY` | --sorted-order `SORTEDORDER`]] [[--filter-column `FILTERCOLUMN` | --filter `FILTER`]]

Where *SORTEDBY* is one of the table column names returned for the selected command target, and *SORTEDORDER* can be either ASC for an ascending sort, or DESC for a descending sort. See Section 3.1.3.2, “Sorting” for more information.

Where *FILTERCOLUMN* is one of the table column names returned for the selected command target, and *FILTER* is the text that you wish to match to perform your filtering. See Section 3.1.3.3, “Filtering” for more information.

Description

Use the `list` command to obtain tabulated listings of information about different components or activities within the Oracle PCA. The `list` command can frequently be used to obtain identifiers that can be used in conjunction with many other commands to perform various actions or to obtain more detailed information about a specific component or task. The `list` command also supports sorting and filtering capabilities to allow you to order information or to limit information so that you are able to identify specific items of interest quickly and easily.

Options

The following table shows the available options for this command.

Option	Description
{ `compute-node` | `config-error` | `management-node` | `network-cloud` | `opus-port` | `storage-cloud` | `task` | `wwpn-info` }	The command target to list information for.
[--json]	Return the output of the command in JSON format
[--sorted-by `SORTEDBY`]	Sort the table by the values within a particular column in the table, specified by replacing `SORTEDBY` with the name of the column that should be used to perform the sort.
[--sorted-order `SORTEDORDER`]	Used to specify the sort order, which can either be `ASC` for an ascending sort, or `DESC` for a descending sort. You must use the `--sorted-by` option in conjunction with this option.
[--filter-column `FILTERCOLUMN`]	Filter the table for a value within a particular column in the table, specified by replacing `FILTERCOLUMN` with the name of the column that should be used to perform the sort. You must use the `--filter` option in conjunction with this option.
[--filter `FILTER`]	The filter that should be applied to values within the column specified by the `--filter-column` option.

Note

When the target of the `list` command is either `management-node` or `compute-node`, the Node_State column is used to show the software completeness of a node through the provisioning process. It does not indicate the system status of a node.

Examples

Example 3.8 List all management nodes

```
PCA> list management-node
```

```
Management_Node IP_Address    Node_State    ILOM_MAC            ILOM_State    Master
--------------- ----------    ----------    --------            ----------    ------
ovcamn06r1      192.168.4.4   RUNNING       00:10:e0:0f:5b:59   running
ovcamn05r1      192.168.4.3   RUNNING       00:10:e0:0f:8e:fb   running       Yes
----------------
2 rows displayed

Status: Success
```

Example 3.9 List all compute nodes

```
PCA> list compute-node

Compute_Node  IP_Address    Node_State  ILOM_MAC            ILOM_State
------------  ----------    ----------  --------            ----------
ovcacn13r1    192.168.4.30  DEAD        00:10:e0:46:9c:59   initializing_stage_wait_for_ovm
ovcacn14r1    192.168.4.5   RUNNING     00:10:e0:46:9e:45   running
ovcacn11r1    192.168.4.31  DEAD        00:10:e0:0f:57:93   initializing_stage_wait_for_ovm
----------------
3 rows displayed

Status: Success
```

Example 3.10 List Network Clouds

```
PCA> list network-cloud

Cloud_Name                          Trunkmode    Description
----------                          ---------    -----------
vm_public_vlan                      True         2014-07-12 02:11:27
mgmt_public_eth                     True         2014-07-12 02:09:42
mgmt_pvi                            False        2014-07-12 02:08:38
vm_private                          False        2014-07-12 02:10:35
discovered-network-cloud            False
----------------
5 rows displayed

Status: Success
```

Example 3.11 List ports on each Oracle Switch ES1-24 using a Filter

Note that the CLI uses the internal alias for an Oracle Switch ES1-24, referencing an 'opus' switch or port.

```
PCA> list opus-port --filter-column=Dest_Port --filter=Net\ Mgt

Dest     Dest_Port  Hostname        Key        OPUS     RACK  RU   Src_Port  Type
----     ---------  --------        ---        ----     ----  --   --------  ----
22       Net Mgt    ovcasw22r1      OPUS-1-9   OPUS-1   1     22   9         switch
19       Net Mgt    ovcasw19r1      OPUS-1-10  OPUS-1   1     19   10        switch
PDU-A    Net Mgt    ovcapoPDU-Ar1   OPUS-1-20  OPUS-1   1     0    20        power
15       Net Mgt    ovcasw15r1      OPUS-2-9   OPUS-2   1     15   9         switch
20       Net Mgt    ovcasw20r1      OPUS-2-10  OPUS-2   1     20   10        switch
PDU-B    Net Mgt    ovcapoPDU-Br1   OPUS-2-20  OPUS-2   1     0    20        power
16       Net Mgt    ovcasw16r2      OPUS-3-10  OPUS-3   2     16   10        switch
PDU-A    Net Mgt    ovcapoPDU-Ar2   OPUS-3-20  OPUS-3   2     0    20        power
20       Net Mgt    ovcasw20r2      OPUS-4-10  OPUS-4   2     20   10        switch
PDU-B    Net Mgt    ovcapoPDU-Br2   OPUS-4-20  OPUS-4   2     0    20        power
16       Net Mgt    ovcasw16r3      OPUS-5-10  OPUS-5   3     16   10        switch
PDU-A    Net Mgt    ovcapoPDU-Ar3   OPUS-5-20  OPUS-5   3     0    20        power
20       Net Mgt    ovcasw20r3      OPUS-6-10  OPUS-6   3     20   10        switch
PDU-B    Net Mgt    ovcapoPDU-Br3   OPUS-6-20  OPUS-6   3     0    20        power
-----------------
14 rows displayed
```

```
Status: Success
```

Example 3.12 List all tasks

```
PCA> list task

Task_ID                          Status  Progress Start_Time
-------                          ------  -------- ----------
1cab8b9ed7d74bf2a9f930de888d896b  FAILURE      100 06-01-2014 21:00:02
01980f4ba1df4f669cef7a449fd0f435  FAILURE      100 06-01-2014 21:00:01
-----------------
2 rows displayed

Status: Success
```

Example 3.13 List Storage Clouds For External Fibre Channel Storage

```
PCA> list storage-cloud

Cloud_Name                         Description
----------                         -----------
discovered-storage-cloud
Cloud_D                            Default Storage Cloud ru15 port2 - Do not delete or modify
Cloud_A                            Default Storage Cloud ru22 port1 - Do not delete or modify
Cloud_C                            Default Storage Cloud ru15 port1 - Do not delete or modify
Cloud_B                            Default Storage Cloud ru22 port2 - Do not delete or modify
----------------
4 rows displayed

Status: Success
```

Example 3.14 List WWPNs To Configure External Fibre Channel Storage

```
PCA> list wwpn-info

WWPN                     vHBA     Cloud_Name     Server       Type  Alias
-------------            ----     ----------     ---------    ----- -------------
50:01:39:70:00:4F:91:00  vhba01   Cloud_A        ovcamn05r1   MN    ovcamn05r1-Cloud_A
50:01:39:70:00:4F:91:02  vhba01   Cloud_A        ovcamn06r1   MN    ovcamn06r1-Cloud_A
50:01:39:70:00:4F:91:04  vhba01   Cloud_A        ovcacn07r1   CN    ovcacn07r1-Cloud_A
50:01:39:70:00:4F:91:06  vhba01   Cloud_A        ovcacn08r1   CN    ovcacn08r1-Cloud_A
50:01:39:70:00:4F:91:01  vhba02   Cloud_B        ovcamn05r1   MN    ovcamn05r1-Cloud_B
50:01:39:70:00:4F:91:03  vhba02   Cloud_B        ovcamn06r1   MN    ovcamn06r1-Cloud_B
50:01:39:70:00:4F:91:05  vhba02   Cloud_B        ovcacn07r1   CN    ovcacn07r1-Cloud_B
50:01:39:70:00:4F:91:07  vhba02   Cloud_B        ovcacn08r1   CN    ovcacn08r1-Cloud_B
50:01:39:70:00:4F:F1:06  vhba03   Cloud_C        ovcacn08r1   CN    ovcacn08r1-Cloud_C
50:01:39:70:00:4F:F1:04  vhba03   Cloud_C        ovcacn07r1   CN    ovcacn07r1-Cloud_C
50:01:39:70:00:4F:F1:02  vhba03   Cloud_C        ovcamn06r1   MN    ovcamn06r1-Cloud_C
50:01:39:70:00:4F:F1:00  vhba03   Cloud_C        ovcamn05r1   MN    ovcamn05r1-Cloud_C
50:01:39:70:00:4F:F1:07  vhba04   Cloud_D        ovcacn08r1   CN    ovcacn08r1-Cloud_D
50:01:39:70:00:4F:F1:05  vhba04   Cloud_D        ovcacn07r1   CN    ovcacn07r1-Cloud_D
50:01:39:70:00:4F:F1:03  vhba04   Cloud_D        ovcamn06r1   MN    ovcamn06r1-Cloud_D
50:01:39:70:00:4F:F1:01  vhba04   Cloud_D        ovcamn05r1   MN    ovcamn05r1-Cloud_D
-----------------
16 rows displayed

Status: Success
```

Example 3.15 List All Configuration Errors

```
PCA> list config-error

ID        Module                     Host            Timestamp
--        ------                     ----            ---------
87        Management node password   192.168.4.4     Thu Aug 21 02:45:42 2014
84        MySQL management password  192.168.4.216   Thu Aug 21 02:44:54 2014
```

```
---------------
2 rows displayed

Status: Success
```

3.2.5 reprovision

The `reprovision` command can be used to trigger reprovisioning for a specified compute node within the Oracle Private Cloud Appliance.

Syntax

`reprovision` { `compute-node` } `node` [`--json`] [`--mac` `MAC`] [`--force`]

where `node` is the compute node name for the compute node that should be reprovisioned, and `MAC` is the MAC address of the network interface that the ILOM is available at.

Description

Use the `reprovision` command to reprovision a specified compute node. The provisioning process is described in more detail in Section 1.4, “Provisioning and Orchestration”.

The `reprovision` command triggers a task that is responsible for handling the reprovisioning process and exits immediately with status 'Success' if the task has been successfully generated. This does not mean that the reprovisioning process itself has completed successfully. To monitor the status of the reprovisioning task, you can use the `list task` and `show task` commands. You can also monitor the log file for information relating to provisioning tasks. The location of the log file can be obtained by checking the Log_File parameter when you run the `show system-properties` command. See Example 3.20, “Show System Properties” for more information.

Options

The following table shows the available options for this command.

Option	Description
{ `compute-node` }	The command target to perform the reprovision operation against.
`--json`	Return the output of the command in JSON format.
`--mac` `MAC`	The ILOM MAC address.
`--force`	Force reprovisioning of a compute node in the case where provisioning has already successfully completed. You can only force reprovisioning on a compute node that does not have any virtual machines running.

Examples

Example 3.16 Reprovisioning a Compute Node

```
PCA> reprovision compute-node ovcacn11r1
Status: Success
```

3.2.6 rerun

Triggers a configuration task to re-run on the Oracle Private Cloud Appliance.

Syntax

```
rerun { config-task } id [ --json ]
```

where `id` is the identifier for the configuration task that must be re-run.

Description

Use the `rerun` command to re-initiate a configuration task that has failed. Use the `list config-error` command to view the configuration tasks that have failed and the associated identifier that you should use in conjunction with this command. See Example 3.15, “List All Configuration Errors” for more information.

You can use the `show task` command to view the status of the task after you have triggered it to run. See Example 3.21, “Show Task” for more information.

Options

The following table shows the available options for this command.

Option	Description
{ `config-task` }	The command target to perform the rerun operation against.
`--json`	Return the output of the command in JSON format

Examples

Example 3.17 Re-run a configuration task

```
PCA> rerun config-task 84
Status: Success
```

3.2.7 set

Sets the value for a system property on the Oracle Private Cloud Appliance.

Syntax

```
set { system-property } { ftp_proxy | http_proxy | https_proxy | log_count | log_file |
log_level | log_size } value [ --json ]
```

where `value` is the value for the system property that you are setting.

Description

Use the `set` command to set the value for a system property on the Oracle Private Cloud Appliance.

Important

The `set system-property` command only affects the settings for the management node where it is run. If you change a setting on the active management node, using this command, you should connect to the passive management node and run the equivalent command there as well, to keep the two systems synchronized. This is the only exception where it is necessary to run a CLI command on the passive management node.

You can use the `show system-properties` command to view the values of various system properties at any point. See Example 3.20, “Show System Properties” for more information.

Important

Changes to system-properties usually require that you restart the service for the change to take effect. To do this, you must run `service ovca restart` in the shell of the active management node after you have set the system property value.

Options

The following table shows the available options for this command.

Option	Description
{ `system-property` }	The command target to perform the set operation against.
`--json`	Return the output of the command in JSON format

Examples

Example 3.18 Changing the location of the pca.log file

```
PCA> set system-property log_file /mnt/filer/logs/pca.log
Status: Success
```

Note

Log configuration through the CLI is described in more detail in Section 7.2, “Setting the Oracle Private Cloud Appliance Logging Parameters”.

Example 3.19 Configuring and unconfiguring an HTTP proxy

```
PCA> set system-property http_proxy http://proxy.example.com:8080
Status: Success

PCA> set system-property http_proxy ''
Status: Success
```

Note

Proxy configuration through the CLI is described in more detail in Section 7.1, “Adding Proxy Settings for Oracle Private Cloud Appliance Updates”.

3.2.8 show

The `show` command can be used to view information about particular objects such as tasks, rack layout or system properties. Unlike the `list` command, which applies to a whole target object type, the `show` command displays information specific to a particular target object. Therefore, it is usually run by specifying the command, the target object type and the object identifier.

Syntax

`show` { `cloud-wwpn` | `network-cloud` | `rack-layout` | `storage-cloud` | `system-properties` | `task` | `vhba-info` } `object` [`--json`]

Where `object` is the identifier for the target object that you wish to show information for. The following table provides a mapping of identifiers that should be substituted for `object`, depending on the command target.

Command Target	Object Identifier
cloud-wwpn	Storage Cloud Name

Command Target	Object Identifier
network-cloud	Network Cloud Name
rack-layout	Rack Name
storage-cloud	Storage Cloud Name
task	Task ID
vhba-info	Compute Node Name

Note that you can use tab completion to help you correctly specify the `object` for the different command targets. You do not need to specify an `object` if the command target is `system-properties`.

Description

Use the `show` command to view information specific to a particular target object, identified by specifying the identifier for the object that you wish to view. The exception to this is the option to view `system-properties`, for which no identifier is required.

Frequently, the `show` command may display information that is not available using the `list` command in conjunction with its filtering capabilities.

Options

The following table shows the available options for this command.

Option	Description
{ `cloud-wwpn` | `network-cloud` | `rack-layout` | `storage-cloud` | `system-properties` | `task` | `vhba-info` }	The command target to show information for.
[`--json`]	Return the output of the command in JSON format

Examples

Example 3.20 Show System Properties

Note

This command only displays the system properties for the management node where it is run. If the system properties have become unsynchronized across the two management nodes, the information reflected by this command may not apply to both systems. You can run this command on either the active or passive management node if you need to check that the configurations match.

```
PCA> show system-properties

----------------------------------------
HTTP_Proxy
HTTPS_Proxy
FTP_Proxy
Log_File             /var/log/ovca.log
Log_Level            DEBUG
Log_Size (MB)        50
Log_Count            25
----------------------------------------
```

```
Status: Success
```

Example 3.21 Show Task

```
PCA> show task 8efc92ef13504dacaa3870348d2a54d7

----------------------------------------
Task_Name            backup
Status               SUCCESS
Progress             100
Start_Time           06-27-2014 09:00:01
End_Time             06-27-2014 09:00:01
Pid                  107429
Result               0
----------------------------------------

Status: Success
```

Example 3.22 Show Rack Layout

```
PCA> show rack-layout x3-2_base

Index Static Sub_Type   Type            Units
----- ------ --------   ----            -----
    0      1 pdu1       pdu             [0]
    1      1 head2      storage         [4, 3, 2, 1]
    2      1 head1      storage         [4, 3, 2, 1]
    3      1 jbod       storage         [4, 3, 2, 1]
    4      1 jbod       storage         [4, 3, 2, 1]
    5      1 mgmt2      management      [5]
    6      1 mgmt1      management      [6]
    7      0 None       compute         [7]
    8      0 None       compute         [8]
    9      0 None       compute         [9]
   10      0 None       compute         [10]
   11      0 None       compute         [11]
   12      0 None       compute         [12]
   13      0 None       compute         [13]
   14      0 None       compute         [14]
   15      1 xsigo2     xsigo           [18, 17, 16, 15]
   16      1 xsigo2     xsigo           [18, 17, 16, 15]
   17      1 xsigo2     xsigo           [18, 17, 16, 15]
   18      1 xsigo2     xsigo           [18, 17, 16, 15]
   19      1 ibsw2      ibswitch        [19]
   20      1 ibsw1      ibswitch        [20]
   21      1 opus1      opus            [21]
   22      1 xsigo1     xsigo           [25, 24, 23, 22]
   23      1 xsigo1     xsigo           [25, 24, 23, 22]
   24      1 xsigo1     xsigo           [25, 24, 23, 22]
   25      1 xsigo1     xsigo           [25, 24, 23, 22]
   26      0 None       compute         [26]
   27      0 None       compute         [27]
   28      0 None       compute         [28]
   29      0 None       compute         [29]
   30      0 None       compute         [30]
   31      0 None       compute         [31]
   32      0 None       compute         [32]
   33      0 None       compute         [33]
   34      0 None       compute         [34]
   35      0 None       compute         [35]
   36      0 None       compute         [36]
   37      0 None       compute         [37]
   38      0 None       compute         [38]
   39      0 None       compute         [39]
   40      0 None       compute         [40]
   41      0 None       compute         [41]
   42      0 None       compute         [42]
```

```
-----------------
43 rows displayed

Status: Success
```

Example 3.23 Show the Ports and VNICs in the mgmt_public_eth Network Cloud

```
PCA> show network-cloud mgmt_public_eth

----------------------------------------
Cloud_Name       mgmt_public_eth
Trunkmode        True
Description      2014-07-12 02:09:42
Ports            ovcasw22r1:5:1, ovcasw22r1:4:1, ovcasw15r1:5:1, ovcasw15r1:4:1
Vnics            ovcacn27r1-eth5, ovcacn13r1-eth5, ovcacn14r1-eth5, ovcacn09r1-eth5,
                 ovcacn36r1-eth5, ovcamn06r1-eth5, ovcacn29r1-eth5, ovcacn07r1-eth5,
                 ovcacn12r1-eth5, ovcacn30r1-eth5, ovcacn42r1-eth5, ovcacn40r1-eth5,
                 ovcacn10r1-eth5, ovcacn28r1-eth5, ovcacn41r1-eth5, ovcamn05r1-eth5,
                 ovcacn38r1-eth5, ovcacn34r1-eth5, ovcacn37r1-eth5, ovcacn11r1-eth5,
                 ovcacn26r1-eth5, ovcacn39r1-eth5, ovcacn35r1-eth5
----------------------------------------

Status: Success
```

Example 3.24 Show the Storage Cloud

```
PCA> show storage-cloud Cloud_A

----------------------------------------
Cloud_Name       Cloud_A
Description      Default Storage Cloud ru22 port1 - Do not delete or modify
Ports            ovcasw22r1:12:1, ovcasw22r1:3:1
vHBAs            ovcacn11r1-vhba01, ovcacn29r1-vhba01, ovcacn12r1-vhba01,
                 ovcacn09r1-vhba01, ovcacn14r1-vhba01, ovcacn28r1-vhba01,
                 ovcacn08r1-vhba01, ovcacn30r1-vhba01, ovcacn26r1-vhba01,
                 ovcacn13r1-vhba01, ovcacn27r1-vhba01, ovcacn31r1-vhba01,
                 ovcacn10r1-vhba01, ovcamn06r1-vhba01, ovcamn05r1-vhba01
----------------------------------------

Status: Success
```

Example 3.25 Show the WWPNs for a particular Storage Cloud

```
PCA> show cloud-wwpn Cloud_A

----------------------------------------
Cloud_Name       Cloud_A
WWPN_List        50:01:39:70:00:58:91:1C, 50:01:39:70:00:58:91:1A,
                 50:01:39:70:00:58:91:18, 50:01:39:70:00:58:91:16,
                 50:01:39:70:00:58:91:14, 50:01:39:70:00:58:91:12,
                 50:01:39:70:00:58:91:10, 50:01:39:70:00:58:91:0E,
                 50:01:39:70:00:58:91:0C, 50:01:39:70:00:58:91:0A,
                 50:01:39:70:00:58:91:08, 50:01:39:70:00:58:91:06,
                 50:01:39:70:00:58:91:04, 50:01:39:70:00:58:91:02,
                 50:01:39:70:00:58:91:00
----------------------------------------

Status: Success
```

Example 3.26 Show the vHBA configuration for a Compute Node

```
PCA> show vhba-info ovcacn10r1

vHBA_Name      Cloud        WWNN                     WWPN
-------------  -----------  -------------            -------------
```

```
vhba03          Cloud_C      50:01:39:71:00:58:B1:04   50:01:39:70:00:58:B1:04
vhba02          Cloud_B      50:01:39:71:00:58:91:05   50:01:39:70:00:58:91:05
vhba01          Cloud_A      50:01:39:71:00:58:91:04   50:01:39:70:00:58:91:04
vhba04          Cloud_D      50:01:39:71:00:58:B1:05   50:01:39:70:00:58:B1:05
----------------
4 rows displayed

Status: Success
```

3.2.9 start

The `start` command is non-functional in this release, but is reserved for future use. Attempting to run this command returns an error notifying you that the feature is disabled.

3.2.10 stop

The `stop` command is non-functional in this release, but is reserved for future use. Attempting to run this command returns an error notifying you that the feature is disabled.

Chapter 4 Managing the Oracle VM Virtual Infrastructure

Table of Contents

4.1 Guidelines and Limitations 58
4.2 Logging in to the Oracle VM Manager Web UI 58
4.3 Monitoring Health and Performance in Oracle VM 59
4.4 Creating and Managing Virtual Machines 59
4.5 Managing Virtual Machine Resources 62
4.6 Configuring Network Resources for Virtual Machines 64
4.7 Viewing and Managing Storage Resources 66
4.8 Tagging of Resources in Oracle VM Manager 67
4.9 Managing Jobs and Events 67

Warning

Access to the Oracle VM Manager web user interface and command line interface is provided without restrictions. The configuration of Oracle Private Cloud Appliance (PCA) components within Oracle VM Manager is automatic and handled by the Oracle PCA provisioning process. Altering the configuration of these components directly within Oracle VM Manager is not supported and may result in the malfunction of the appliance.

Use of Oracle VM Manager in the context of Oracle Private Cloud Appliance should be limited to the management and creation of virtual machines.

Configuring additional storage, creating repositories, and setting up additional networks specifically for the use of virtual machines is possible. However, this should be done carefully, to avoid disrupting the configuration specific to the Oracle PCA.

Management of virtual machines and your Oracle VM environment is achieved using the Oracle VM Manager Web UI (User Interface). While Oracle VM Manager does provide a command line interface, use of this on your Oracle PCA should only be attempted by advanced users with a thorough understanding of Oracle VM and the limitations of its usage within an Oracle PCA context.

The information provided in here, is a description of the Oracle VM Manager Web UI within the context of the Oracle PCA. Where particular actions within the Oracle VM Manager Web UI are referenced, a link to the appropriate section within the Oracle VM User's Guide is provided. The complete Oracle VM User's Guide is available at this URL: http://docs.oracle.com/cd/E35328_01/E35332/html/index.html.

Note

When consulting the Oracle VM documentation directly, keep in mind the limitations imposed by using it within Oracle PCA.

The Oracle VM Manager Web UI is available at the virtual IP address that you configured for your management nodes during installation. This virtual IP address is automatically assigned to whichever management node is currently the active node within the cluster. If the management node becomes unavailable, the standby management node is promoted to the active role and takes over the IP address automatically. See Section 1.5, “High Availability” for more information on management node failover.

The Oracle VM Manager Web UI is configured to listen for HTTPS requests on port 7002.

4.1 Guidelines and Limitations

The Oracle VM Manager Web User Interface is provided without any software limitation to its functionality. Once your appliance has been provisioned, the Oracle VM environment is fully configured and ready to use for the deployment and management of your virtual machines. In this section, the operations that are explicitly not permitted, and those that may be attempted with due care, are presented as guidelines and limitations that should be followed when working within Oracle VM Manager.

The following actions must not be performed, except if Oracle gives specific instructions to do so.

Do Not:

- attempt to discover or modify or remove servers or their configuration;
- attempt to add or modify or remove server pools or their configuration;
- attempt to move servers out of the existing server pool;
- attempt to add or modify or remove affinity groups or server processor compatibility groups;
- attempt to modify or remove the existing repository named **Rack1-repository**, or the local disk repositories;
- attempt to delete or modify any of the preconfigured networks;
- attempt to connect virtual machines to the appliance management network;
- attempt to add or configure VLAN Groups, except to add VLAN tags to the default VLAN Group configuration;
- attempt to modify or delete any existing Storage elements that are already configured within Oracle VM;
- attempt to configure global settings in the **Tools and Resources** tab, such as the NTP or YUM Update configuration.

While unlikely to cause a problem with your appliance, the following actions are not recommended:

Not Recommended:

- add or configure additional networks within Oracle VM Manager;
- use the Oracle VM Manager Command Line Interface or Utilities.

If you ignore this advice, the Oracle PCA automation, which uses specific naming conventions to label and manage assets, may fail. Out-of-band configuration changes would not be known to the orchestration software of the Oracle PCA.

4.2 Logging in to the Oracle VM Manager Web UI

To open the Login page of the Oracle VM Manager Web UI, enter the following address in a Web browser:

https://*`manager-vip`*:7002/ovm/console

Where, *`manager-vip`* refers to the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation. By using the virtual IP address, you ensure that you always access the Oracle VM Manager Web UI on the active management node.

Important

You must ensure that if you are accessing Oracle VM Manager through a firewalled connection, the firewall is configured to allow TCP traffic on the port that Oracle VM Manager is using to listen for connections.

Enter your Oracle VM Manager administration user name in the **Username** field. This is the administration user name you configured during installation. Enter the password for the Oracle VM Manager administration user name in the **Password** field.

Important

The Oracle VM Manager Web UI makes use of cookies in order to store session data. Therefore, to successfully log in and use the Oracle VM Manager Web UI your web browser must accept cookies from the Oracle VM Manager host.

4.3 Monitoring Health and Performance in Oracle VM

The **Health** tab provides a view of the health of the compute nodes and the server pool within your environment. This information complements the Hardware View provided in the Oracle PCA Dashboard. See Section 2.2, “Hardware View” for more information.

The **Statistics** subtab available on the Health tab provides statistical information, including real-time graphs, for CPU and memory usage. These statistics can be viewed at a global scale to determine overall usage, per server to see the performance of each individual compute node, or per virtual machine to help track the usage and resource requirements for any of the virtual machines within your environment.

For detailed information on using the Health tab, please refer to the section entitled Health Tab in the Oracle VM User's Guide.

4.4 Creating and Managing Virtual Machines

The **Servers and VMs** tab is used to create and manage your virtual machines. By default, compute nodes are listed as belonging to a single server pool called **Rack1_ServerPool**. The configuration of this server pool must not be altered. There is no need to discover servers, as compute nodes are automatically provisioned and discovered within an Oracle PCA. Editing the configuration of the server pool, servers, affinity groups and processor compatibility groups is not supported. The primary purpose of this tab within the Oracle PCA context is to create and manage your virtual machines.

Virtual machines can be created using:

- ISO files in a repository (hardware virtualized only)
- Mounted ISO files on an NFS, HTTP or FTP server (paravirtualized only)
- Virtual machine templates (by cloning a template)
- Existing virtual machines (by cloning a virtual machine)
- Virtual machine assemblies

Virtual machines require most installation resources to be located in the storage repository, managed by Oracle VM Manager, with the exception of mounted ISO files for paravirtualized guests. See Section 4.5, “Managing Virtual Machine Resources” for more information on importing these resources into the Oracle PCA repository.

Before you create a virtual machine that requires network connectivity, or a paravirtualized machine which requires network connectivity to perform the operating system install, you must generate virtual

network interfaces using the VNIC Manager. See Section 4.6, “Configuring Network Resources for Virtual Machines” for information on using the VNIC Manager.

The following list provides an outline of actions that you can perform in this tab, with links to the relevant documentation within the Oracle VM User's Guide:

Managing Virtual Machines

- Create a virtual machine

  You can create a virtual machine following the instructions provided in the section entitled Creating a Virtual Machine.

  You do not need to create any additional server pools. You need only ensure that your installation media has been correctly imported into the Oracle PCA repository.

- View virtual machine information and events

  You can view information about your virtual machine or access virtual machine events by following the information outlined in the section entitled Viewing Virtual Machine Information and Events.

- Edit a virtual machine

  You can edit virtual machine parameters as described in the section entitled Editing a Virtual Machine.

  You can also convert a virtual machine to use paravirtualization. This involves making changes to the virtual machine itself, and then editing its Oracle VM Domain Type. Refer to the section entitled Converting to Paravirtualized Guests or Installing Paravirtualized Drivers for information on converting a virtual machine to use paravirtualization.

- Start a virtual machine

  Further information is provided in the section entitled Starting Virtual Machines.

- Stop a virtual machine

  Further information is provided in the section entitled Stopping (Shutting Down) Virtual Machines.

- Kill a virtual machine

  Further information is provided in the section entitled Killing Virtual Machines.

- Restart a virtual machine

  Further information is provided in the section entitled Restarting Virtual Machines.

- Suspend a virtual machine

  Further information is provided in the section entitled Suspending Virtual Machines.

- Resume a virtual machine

  Further information is provided in the section entitled Resuming a Virtual Machine.

- Move a virtual machine between repositories

  It is possible to create alternate repositories if you have extended the system with external storage. If you have an additional repository, this function can be used to move a virtual machine from one

repository to another. Refer to the section entitled Moving Virtual Machines Between Repositories for more information.

- Move a virtual machine from one server to another

  Further information is provided in the section entitled Moving Virtual Machines Between Oracle VM Servers.

- Move a virtual machine to or from the Unassigned Virtual Machines folder

  Further information is provided in the section entitled Moving Virtual Machines To/From Unassigned Virtual Machines Folder.

- Migrate a virtual machine

  Because there is only a single server pool available in an Oracle PCA base rack, migration of virtual machines can only be achieved between servers and between a server and the Unassigned Virtual Machines folder. Modifying Server Processor Compatibility Groups is not permitted.

  Note

  Compute nodes of different hardware generations operate within the same server pool but belong to different CPU compatibility groups. Since live migration between CPU compatibility groups is not supported, virtual machines have to be cold-migrated between compute nodes of different generations. For more information about CPU compatibility groups, please refer to the section entitle Server Processor Compatibility Groups.

  Information on migrating virtual machines is provided in the section entitled Migrating Virtual Machines.

- Delete a virtual machine

  Further information is provided in the section entitled Deleting Virtual Machines.

- Send a message to a virtual machine

  If you have installed Oracle VM Guest Additions within your virtual machine, you can use the Oracle VM Messaging framework to send messages to your virtual machines to trigger actions within a virtual machine. Refer to the section entitled Sending Messages to Virtual Machines for more information.

- Connect to a virtual machine console

  Further information is provided in the section entitled Connecting to a Virtual Machine.

Monitoring Compute Node Information and Events

- View Oracle VM Server information and events

  For more information on the options available to you, refer to the section entitled Viewing Oracle VM Server Information and Events.

Figure 4.1 A view of the Servers and VMs tab

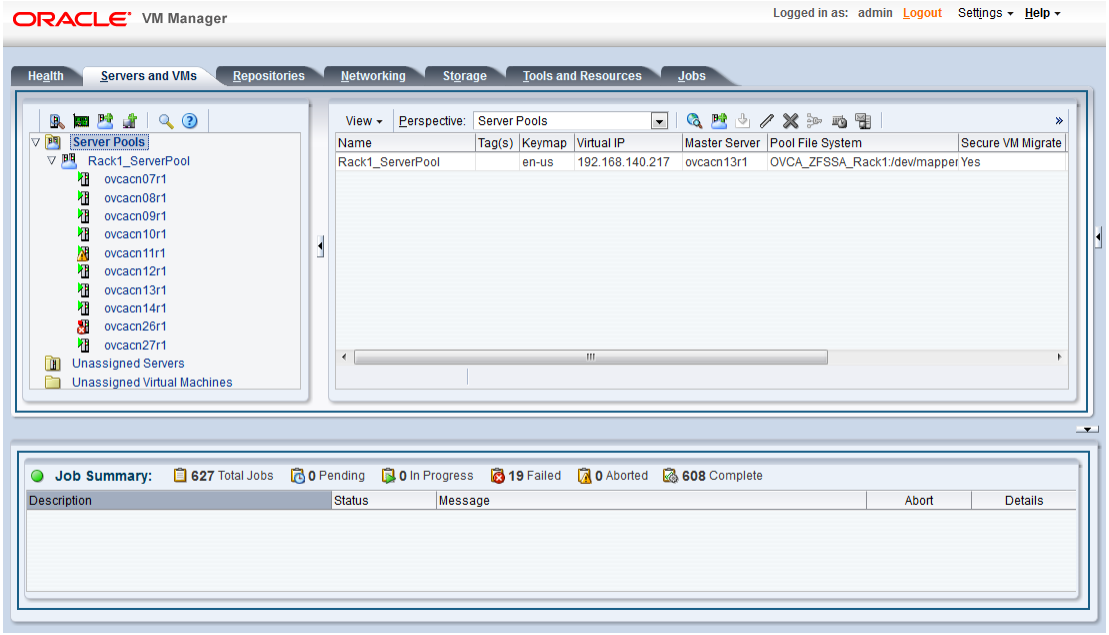

4.5 Managing Virtual Machine Resources

The **Repositories** tab provides a view of the Oracle PCA repository. By default, a shared repository is configured on the ZFS storage appliance and named **Rack1-repository**. Additional local repositories are configured using the free disk space of each compute node. None of the default repository configurations may be altered.

Caution

Using local storage on the compute nodes has implications that you should take into account when planning the deployment of your virtual environment. For example:

- Virtual machines with resources in a local storage repository cannot be migrated to another compute node.
- Templates, assemblies and ISOs in local storage repositories cannot be used to create virtual machines on another compute node.
- If a compute node becomes unavailable, its locally stored virtual machines and resources cannot be restored or migrated to another compute node for continued service.
- The virtual machines and resources in local storage repositories are not protected by automatic failover and high-availability mechanisms normally offered by a clustered Oracle VM server pool with shared storage repository.

The Repositories tab is used to manage virtual machine resources, such as installation media and virtual disks. From this tab, it is possible to create, import or clone Oracle VM templates, assemblies and ISO image files. It is also possible to create, modify, or clone virtual disks here. The following list provides an outline of actions that you can perform in this tab, with links to the relevant documentation within the Oracle VM User's Guide:

- Manage Virtual Machine Templates
  - Import a template
  - Create a template
  - Edit a template
  - Delete a template
  - Clone a template
  - Move a template
  - Manage a template clone customizer

  All documentation for these actions can be found in the section entitled Virtual Machine Templates.
- Manage Virtual Machine Assemblies
  - Import an assembly
  - Create a template from an assembly
  - Edit an assembly
  - Delete an assembly
  - Refresh an assembly

  All documentation for these actions can be found in the section entitled Assemblies.
- Manage Virtual Machine ISO Image Files
  - Import an ISO
  - Edit an ISO
  - Delete an ISO
  - Clone an ISO

  All documentation for these actions can be found in the section entitled ISO Files (CD/DVD Images).
- Manage Virtual Disks
  - Create a virtual disk
  - Import a virtual disk
  - Edit a virtual disk
  - Delete a virtual disk
  - Clone a virtual disk

  All documentation for these actions can be found in the section entitled Virtual Disks.
- View Virtual Machine Configuration Entries

For more information, refer to the section entitled Virtual Machine Configuration Files.

While it is possible to create additional repositories using alternate storage, this is beyond the scope of a normal Oracle PCA configuration and is not recommended.

4.6 Configuring Network Resources for Virtual Machines

The **Networking** tab is used to manage networks within the Oracle VM environment running on the Oracle PCA.

Caution

By default, a number of networks are defined during factory installation. These **must not be altered** as they are required for the correct operation of the Oracle PCA software layer.

The default networks are set up as follows:

- `192.168.140.0` : the management network

  This is a private network used exclusively for Oracle VM management traffic. Both management nodes and all compute nodes are connected to this network through their `bond0` interface.

- `192.168.40.0` : the storage network

  This is a private IPoIB network used exclusively for traffic to and from the ZFS storage appliance. Both management nodes and both storage controllers are connected to this network through their `bond1` interface.

Additionally, three networks are listed with the **VM Network** role:

- `vm_public_vlan`

  This default network is the standard choice for virtual machines requiring external network connectivity. It supports both tagged and untagged traffic. For untagged traffic it uses VLAN 1 from the VLAN Group named `vm_public_vlan_vgrp`, meaning no additional configuration is required.

  If you prefer to use VLANs for your VM networking, configure the VLAN tags of your choice as follows:

  1. Add the tags or IDs to the VLAN Group `vm_public_vlan_vgrp`.

     The process for editing VLAN Groups is described in the Oracle VM User's Guide in the section entitled Editing a VLAN Group.

  2. Create a new VLAN-only network with the *VM role* for each VLAN tag you added to the VLAN Group. Each new network should be configured like the `vm_public_vlan` network, but with a different VLAN segment.

     The process for creating VLAN-only networks is described in the Oracle VM User's Guide in the section entitled Creating a Network.

  3. Configure your data center network accordingly.

     For details, see Section 7.3, “Configuring Data Center Switches for VLAN Traffic”.

- `vm_private`

This default network is intended for virtual machines requiring network connectivity to other virtual machines hosted on the appliance, but *not* external to the appliance. It uses VLAN 1 from the VLAN Group named `vm_private_vgrp`. To use the VLANs of your choice configure them as follows:

1. Add the tags or IDs to the VLAN Group `vm_private_vgrp`.

   The process for editing VLAN Groups is described in the Oracle VM User's Guide in the section entitled Editing a VLAN Group.

2. Create a new VLAN-only network with the *VM role* for each VLAN tag you added to the VLAN Group. Each new network should be configured like the `vm_private` network, but with a different VLAN segment.

   The process for creating VLAN-only networks is described in the Oracle VM User's Guide in the section entitled Creating a Network.

- `mgmt_public_eth`

  This network is automatically created during the initial configuration of the appliance. It uses the public network that you configured in the Oracle PCA Dashboard. The primary function of this network is to provide access to the management nodes from the data center network, and enable the management nodes to run a number of system services. As long as you have not configured this network with a VLAN tag, it may also be used to provide external untagged network access to virtual machines. The subnet associated with this network is the same as your data center network.

Caution

Always use the `vm_public_vlan` network as your first VM network option. The `mgmt_public_eth` is unavailable for VM networking when configured with a management VLAN. When no management VLAN is configured, it is restricted to untagged VM traffic, and should only be considered if the circumstances require it.

For more information about Oracle PCA network configuration, see Section 1.2.4, “Network Infrastructure”.

Caution

Do not create a new network in Oracle VM using the `eth0` network ports of the servers in the server pool. Those ports are assigned an IP address by the DHCP server on the internal appliance management network (`192.168.4.0`). If virtual machines are connected to this network, they are likely to cause IP conflicts and security issues.

Do not edit or delete any of the networks listed here. Doing so may cause your appliance to malfunction. In an Oracle PCA context, use the Networking tab to configure and manage Virtual NICs and VLANs for use by your virtual machines.

The creation and management of Virtual NICs is important for the facilitation of networking within your virtual machines. Use the **Virtual NICs** subtab to create batches of Virtual NICs that virtual machines are able to use. The process for creating Virtual NICs is described in the Oracle VM User's Guide in the section entitled Managing VNICs.

Figure 4.2 A view of the Networking tab

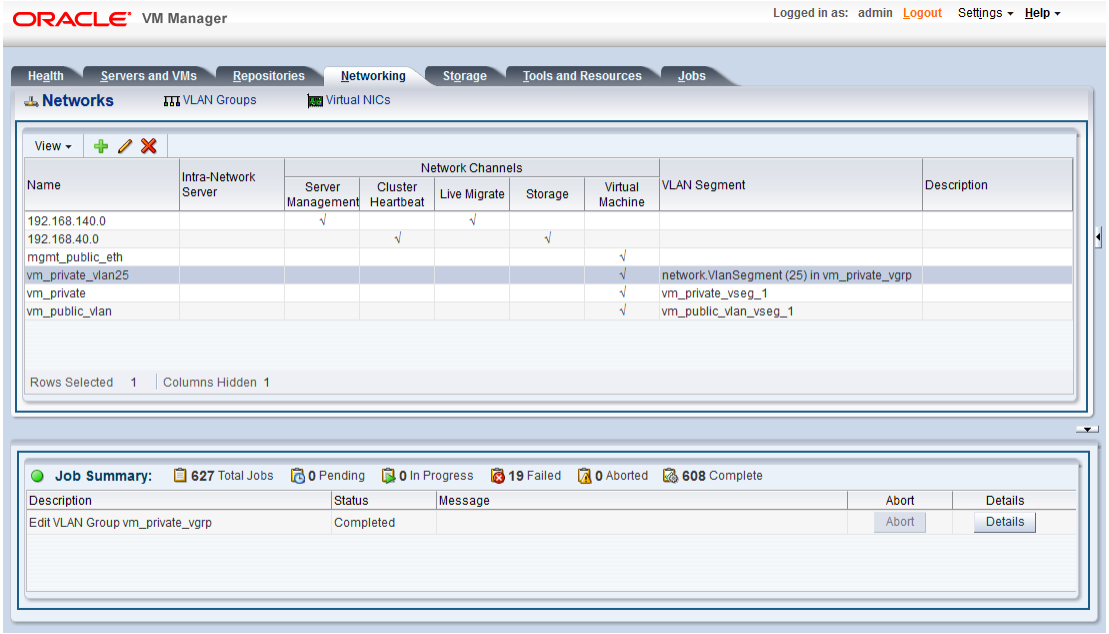

4.7 Viewing and Managing Storage Resources

The storage resources underlying the Oracle PCA repository and the server pool clustering file system are listed under the **Storage** tab within Oracle VM Manager. The ZFS storage is exported through NFS and is listed under the File Servers folder. Do not modify or attempt to delete this storage.

It is possible to attach additional storage for use within Oracle VM, if you intend to create an additional repository to store virtual machine resources. This activity is beyond the scope of a normal Oracle PCA configuration and is not recommended. Information on adding File Servers to your Oracle VM environment is provided in the Oracle VM User's Guide in the section entitled Discovering File Servers.

Note that only NFS storage is supported within Oracle PCA Release 2.0.

You are fully capable of using other networked storage, available on the public network, within your own Virtual Machines. The limitations described here apply only to storage used by the Oracle VM environment running within your Oracle PCA.

Figure 4.3 A view of the Storage tab

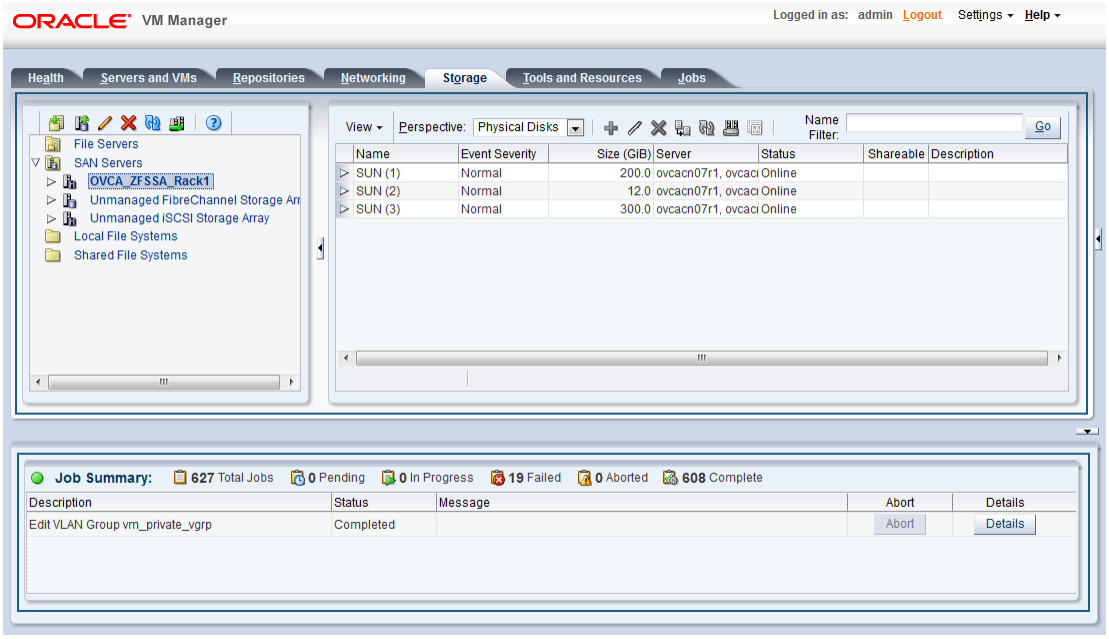

4.8 Tagging of Resources in Oracle VM Manager

The **Tools and Resources** tab is used to configure global settings for Oracle VM and to manage tags, which can be used to identify and group resources. Since many of the global settings such as server update management and NTP configuration are managed automatically within Oracle PCA, you do not need to edit any settings here. Those configuration changes could cause the appliance to malfunction.

You are able to create, edit and delete tags, by following the instructions in the section entitled Tags.

4.9 Managing Jobs and Events

The **Jobs** tab provides a view of the job history within Oracle VM Manager. It is used to track and audit jobs and to help troubleshoot issues within the Oracle VM environment. Jobs and events are described in detail within the Oracle VM User's Guide in the section entitled Working with the Jobs Framework.

Since the Recurring Jobs, described in the Oracle VM User's Guide, are all automated and handled directly by the Oracle PCA, you must not edit any of the settings for recurring jobs.

Chapter 5 Automating Oracle Private Cloud Appliance Service Requests

Table of Contents

5.1 Understanding Oracle Auto Service Request (ASR) 69
5.2 ASR Prerequisites 70
5.3 Installing ASR Manager 71
5.4 Activating ASR Assets 75
5.4.1 Activating Compute Nodes 75
5.4.2 Activating InfiniBand Switches 76
5.4.3 Activating ASR on the ZFS Storage Appliance 78

Oracle Private Cloud Appliance (PCA) is qualified for Oracle Auto Service Request (ASR). ASR is a software feature for support purposes. It is integrated with My Oracle Support and helps resolve problems faster by automatically opening service requests when specific hardware failures occur. Using ASR is optional: the components are available in the Release 2.0 software image but they must be installed and configured in order to enable ASR for your appliance.

For the latest information about installing Oracle Auto Service Request (ASR) on Oracle PCA, refer to the support note with Doc ID 1645856.1.

Caution

Oracle Auto Service Request (ASR) must be installed by an **authorized Oracle Field Engineer**. Request installation of ASR at the time of system install. Installation at a later date will be a Time and Materials charge.

Oracle is continuously analyzing and improving the ASR fault rules to enhance the Oracle support experience. This includes adding, modifying and removing rules to focus on actionable events from ASR assets while filtering non-actionable events. For up-to-date fault coverage details, please refer to the Oracle Auto Service Request documentation page: http://www.oracle.com/technetwork/systems/asr/documentation/index.html.

5.1 Understanding Oracle Auto Service Request (ASR)

ASR is designed to automatically open service requests when specific Oracle PCA hardware faults occur. To enable this feature, the Oracle PCA components must be configured to send hardware fault telemetry to the ASR Manager software. ASR Manager must be installed on the master management node, which needs an active outbound Internet connection using HTTPS or an HTTPS proxy.

When a hardware problem is detected, ASR Manager submits a service request to Oracle Support Services. In many cases, Oracle Support Services can begin work on resolving the issue before the administrator is even aware the problem exists.

ASR detects faults in the most common hardware components, such as disks, fans, and power supplies, and automatically opens a service request when a fault occurs. ASR does not detect all possible hardware faults, and it is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts, within the customer data center. It is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.

An email message is sent to both the My Oracle Support email account and the technical contact for Oracle Private Cloud Appliance to notify them of the creation of the service request. A service request may not be filed automatically on some occasions. This can happen because of the unreliable nature of the SNMP protocol or a loss of connectivity to ASR Manager. Oracle recommends that customers continue to monitor their systems for faults and call Oracle Support Services if they do not receive notice that a service request has been filed automatically.

For more information about ASR, consult the following resources:

- Oracle Auto Service Request web page: http://www.oracle.com/technetwork/systems/asr/overview/index.html.
- Oracle Auto Service Request user documentation: http://docs.oracle.com/cd/E37710_01/index.htm.

5.2 ASR Prerequisites

Before you install ASR, make sure that the prerequisites in this section are met.

Verifying ASR Prerequisites

1. Make sure that you have a valid My Oracle Support account.

   If necessary, create an account at https://support.oracle.com.

2. Ensure that the following are set up correctly in My Oracle Support:

   - technical contact person at the customer site who is responsible for Oracle PCA
   - valid shipping address at the customer site where the Oracle PCA is located, so that parts are delivered to the site where they must be installed

3. Make sure that Oracle Java - JDK 7 (1.7.0_13 or later) is installed on both management nodes in your Oracle PCA. Check the version installed on the system by entering the following command at the Oracle Linux prompt: `java -version`.

   Note

   OpenJDK is not supported by ASR.

   If the installed version does not comply with the ASR prerequisites, download and install the version of JDK included in the Oracle PCA Release 2.0 software image. Unpack the Oracle Java - JDK 7 archive in `/opt/`. Repeat this on the other management node.

```
# cd /opt/
# tar zxvf /nfs/shared_storage/mgmt_image/Server/jdk-7u51-linux-x64.tar.gz
[...]
```

   Note

   Alternatively, you can download the latest version from the Java SE Downloads page: http://www.oracle.com/technetwork/java/javase/downloads/.

   Caution

   Each ASR version is configured to use a particular version of Java. Verify the Java prerequisites whenever you install a new version of ASR. Do not set the newly installed JDK as the system default Java, because that may cause other components to malfunction.

4. Verify connectivity to the Internet using HTTPS.

   For example, try `curl` to test whether you can access https://support.oracle.com.

5.3 Installing ASR Manager

The necessary packages for ASR Manager are included in the Oracle PCA Release 2.0 software image. For ASR Manager to work on Oracle PCA, it must be installed on both management nodes, and failover must be configured so that the ASR Manager role is always fulfilled by the management node that also has the master role.

Installing ASR Components on the Management Nodes

1. Using SSH and an account with superuser privileges, log into the master management node.

   Note

   The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Go to the `Packages` directory of the mounted management software image. Verify that the ASR packages are present.

```
# cd /nfs/shared_storage/mgmt_image/Packages/
# ls -al | grep SUNW
-rw-r--r--  1 root root 12406356 Apr  7 04:37 SUNWsasm-1.5.0-112.rpm
-rw-r--r--  1 root root 24090626 Apr  7 04:37 SUNWswasr-4.7-20140116133106.rpm
```

   Note

   You may also download the latest ASR packages from the Oracle Auto Service Request downloads page: http://www.oracle.com/technetwork/systems/asr/downloads/index.html.

3. Install the ASR Manager packages without verifying dependencies.

```
# rpm -ivh SUNWsasm-1.5.0-112.rpm --nodeps
Copyright 2008,2013 Oracle and/or its affiliates. All rights reserved.

License and Terms of Use for this software are described at
https://support.oracle.com/  (see Terms of Use)

# rpm -ivh SUNWswasr-4.7-20140116133106.rpm --nodeps
Copyright [2008,2012], Oracle and/or its affiliates. All rights reserved.

License and Terms of Use for this software are described at
https://support.oracle.com/ (see Legal Notices and Terms of Use).

Directory /var/opt/SUNWsasm/configuration/caseinfo created.
Directory /var/opt/SUNWsasm/configuration/supportfile created.

ASR Manager Auto Update functionality has been enabled by default.
Please ensure that ASR Manager is registered with ASR backend to get the
software updates.

Warning: rpm-build  package is not installed on this server. ASR Manager
Auto Update functionality will not work unless the rpm-build package is
installed. Auto Update functionality will be disabled until rpm-build
```

```
package is installed. Please install the rpm-build package and then enable
Auto Update by running "asrenable_autoupdate".
Installation of SUNWswasr was successful.
```

Warning

For Oracle PCA the auto-update feature of ASR must not be used.

At the end of the installation, the ASR Manager service (`sasm`) is started automatically.

4. Stop the ASR Manager service and prevent it from starting automatically.

```
# service sasm status
Oracle Automated Service Manager (pid 450596) is running...
# service sasm stop
Stopping Oracle Automated Service Manager...
Stopped.
# chkconfig sasm off
```

5. Configure the `init` service to log when the ASR Manager service starts and stops.

   a. Open the file `/etc/init.d/sasm` for editing.

   b. In the start and stop sections, add the echo commands as shown.

```
'start')
  /opt/SUNWsasm/bin/sasm start-instance
  echo "START: `hostname`  `date "+%D %T"`" >> /nfs/shared_storage/ASRM/ASRHA.log
    ;;

'stop')
  /opt/SUNWsasm/bin/sasm stop-instance
  echo "STOP:  `hostname`  `date "+%D %T"`" >> /nfs/shared_storage/ASRM/ASRHA.log
    ;;
```

   c. Save and close the file `/etc/init.d/sasm`.

6. Make sure that the `sasm` configuration file points to the correct JDK.

   a. Verify and copy the path to the appropriate java executable. If you installed Oracle Java - JDK 7 by following the instructions in the Prerequisites section, it should be `/opt/jdk1.7.0_51/bin/java`.

   b. Open the `sasm` configuration file `/var/opt/SUNWsasm/configuration/config.ini` for editing.

   c. Set the `java.exec` value to the path you copied.

```
# OASM version
#
sasm.version=1.5.0
#
# The JVM lookup-order of OASM has been defined as follows:
# 1. JAVA_HOME/bin
# 2. java.exec value defined below
# 3. java in path
#
# Therefore, note that JAVA_HOME/bin/java takes precedence (if exists) over th
# so remember to unset JAVA_HOME if you prefer to use following java.exec spe
#
# java.exec=/usr/java/default/bin/java
java.exec=/opt/jdk1.7.0_51/bin/java
#
```

```
[...]
```

d. Save and close the file `/var/opt/SUNWsasm/configuration/config.ini`.

7. Preserve the data in the original ASR directories.

```
# mv /var/opt/SUNWsasm /var/opt/orig_SUNWsasm
```

8. Repeat this procedure on the other management node.

Caution

The secondary management node does not need to take over the master role during the ASR installation.

The ASR Manager service must remain stopped on both management nodes.

Configuring ASR Manager Failover

1. Using SSH and an account with superuser privileges, log into the master management node.

Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Create the following directories and log file on the shared storage:

```
# mkdir /nfs/shared_storage/ASRM
# mkdir /nfs/shared_storage/ASRM/VarOptSUNWsasm
# touch /nfs/shared_storage/ASRM/ASRHA.log
```

3. Copy the original ASR data from the master management node local file system to the shared storage.

```
# cp -r /var/opt/orig_SUNWsasm/* /nfs/shared_storage/ASRM/VarOptSUNWsasm
```

4. Create a symlink to the directory on the shared storage.

```
# ln -s /nfs/shared_storage/ASRM/VarOptSUNWsasm /var/opt/SUNWsasm
```

5. Add the `sasm` service to the Oracle PCA system configuration.

a. Open the file `/var/lib/ovca/ovca-system.conf` for editing.

b. Insert a new line to add `sasm` to the foundational services.

```
[mgmt_init]
[...]
foundational_services: ovmm_mysql
                       ovmm
                       tinyproxy
                       xms
                       dhcpd
                       sasm
[...]
```

c. Save and close the file `/var/lib/ovca/ovca-system.conf`.

6. Log into the other management node and execute steps 4 and 5.

a. Create a symlink to the directory on the shared storage.

b. Add the `sasm` service to the Oracle PCA system configuration.

Starting ASR Manager

1. Using SSH and an account with superuser privileges, log into the master management node.

Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Start the ASR Manager service.

Note

You can monitor the process by tailing the log file `/nfs/shared_storage/ASRM/ASRHA.log`.

```
# service sasm start
```

3. Set the ASR Manager log level to *fine*.

```
# /opt/SUNWswasr/bin/asr

asr> get_loglevel
INFO

asr> set_loglevel fine
log.level is set to FINE
Stopping ASR
ASR stopped
Starting ASR
ASR started

asr> get_loglevel
FINE
```

4. Register the ASR Manager.

   ASR Manager (ASRM) can be registered as a stand-alone ASRM or as a relay to another ASRM in your network. Even if other systems at your site already use an ASRM, you can choose to register the Oracle PCA ASRM as stand-alone. This means it communicates directly with the Oracle backend systems, which is the standard registration method.

   For details and instructions, refer to the section entitled "Register the ASR Manager" in the Oracle Auto Service Request Installation and Operations Guide.

   Due to networking restrictions, the Oracle PCA ASRM can only be routed through an existing one. ASRM relay is not possible in the opposite direction. To register the ASRM as a relay, you must enable the HTTP receiver on the existing ASRM, and register the Oracle PCA ASRM to the HTTP receiver web address.

   For details and instructions, refer to the section entitled "Enable HTTP Receiver for ASR Manager Relay, Solaris 11, and VOP" in the Oracle Auto Service Request Installation and Operations Guide.

5. Test ASR Manager failover by rebooting the current master management node.

Note

You can monitor the process by tailing the log file `/nfs/shared_storage/ASRM/ASRHA.log`.

5.4 Activating ASR Assets

The Oracle Private Cloud Appliance components that are qualified as ASR assets are:

- Server compute nodes
- NM2-36P Sun Datacenter InfiniBand Expansion Switches
- ZFS Storage Appliance

Note

There is no ASR support for the Oracle Switch ES1-24 switches and Oracle Fabric Interconnect F1-15s in Oracle Private Cloud Appliance Release 2.0. The two management nodes also must not be activated.

This section provides the instructions to activate each type of Oracle Private Cloud Appliance ASR asset.

5.4.1 Activating Compute Nodes

The ASR activation mechanism requires operations in two separate locations. First the compute node ILOMs are configured to send SNMP traps to the ASR Manager when a failure occurs. Then the ASR Manager is configured to recognize the ILOMs as assets and accept their input. Follow the procedure below to activate the compute nodes as ASR assets.

Activating Compute Nodes as ASR Assets

1. Using SSH, log into the compute node ILOM as `root`.
2. Go to the alert rules directory and display rule set 1.

```
-> cd /SP/alertmgmt/rules
/SP/alertmgmt/rules

-> show 1
  /SP/alertmgmt/rules/1
      Targets:
      Properties:
          type = snmptrap
          level = disable
          destination = 0.0.0.0
          destination_port = 0
          community_or_username = public
          snmp_version = 1
          testrule = (Cannot show property)
->
```

   If the rule set is available (destination = 0.0.0.0), configure it for ASR. Otherwise, look for the next available rule set.

3. Define a new rule set with the following properties:

```
-> cd 1
-> set level=minor destination=192.168.4.216 destination_port=162 snmp_version=2c
```

```
-> show
  /SP/alertmgmt/rules/1
      Targets:
      Properties:
          type = snmptrap
          level = minor
          destination = 192.168.4.216
          destination_port = 162
          community_or_username = public
          snmp_version = 2c
          testrule = (Cannot show property)
      Commands:
          cd
          set
          show
->
```

The *destination* property is a virtual IP in the appliance management network, shared by both management nodes, but active only on the master.

Use a space character to separate rule set properties.

4. Activate the compute node as an ASR asset in ASR Manager.

   a. Using SSH, log into the master management node.

   b. Activate the asset by adding its ILOM IP to the ASR configuration with the following command:

   ```
   # /opt/SUNWswasr/bin/asr activate_asset –i <asset_ilom_ip>
   ```

   Caution

   Verify that you are activating the asset with the IP address of its ILOM, and not the IP assigned to another network interface.

5. Repeat this procedure for each compute node you wish to activate as an ASR asset.

5.4.2 Activating InfiniBand Switches

The ASR activation mechanism requires operations in two separate locations. First the switches' ILOMs are configured to send SNMP traps to the ASR Manager when a failure occurs. Then the ASR Manager is configured to recognize the ILOMs as assets and accept their input. Follow the procedure below to activate the NM2-36P Sun Datacenter InfiniBand Expansion Switches as ASR assets.

Activating InfiniBand Switches as ASR Assets

1. Using SSH, log into the NM2-36P Sun Datacenter InfiniBand Expansion Switch ILOM as `root`.

2. Launch the Service Processor shell.

   ```
   [root@ilom-ovcasw19r1 ~]# spsh
   ->
   ```

3. Go to the alert rules directory and display rule set 1.

   ```
   -> cd /SP/alertmgmt/rules
   /SP/alertmgmt/rules

   -> show 1
     /SP/alertmgmt/rules/1
         Targets:
         Properties:
             type = snmptrap
   ```

```
          level = disable
          destination = 0.0.0.0
          destination_port = 0
          community_or_username = public
          snmp_version = 1
          testrule = (Cannot show property)
      Commands:
          cd
          set
          show
->
```

4. Define a new rule set with the following properties:

```
-> cd 1
-> set level=minor destination=192.168.4.216 destination_port=162 snmp_version=2c
Set 'level' to 'minor'
Set 'destination' to '192.168.4.216'
Set 'destination_port' to '162'
Set 'snmp_version' to '2c'

-> show
  /SP/alertmgmt/rules/1
      Targets:
      Properties:
          type = snmptrap
          level = minor
          destination = 192.168.4.216
          destination_port = 162
          community_or_username = public
          snmp_version = 2c
          testrule = (Cannot show property)
      Commands:
          cd
          set
          show
->
```

   The *destination* property is a virtual IP in the appliance management network, shared by both management nodes, but active only on the master.

   Rule set properties are space-separated.

5. Make sure that SNMP version 2c is enabled on the switch.

```
-> cd /SP/services/snmp
/SP/services/snmp
-> show
  /SP/services/snmp
      Targets:
          communities
          mibs
          users
      Properties:
          engineid = (none)
          port = 161
          servicestate = enabled
          sets = disabled
          v1 = disabled
          v2c = disabled
          v3 = enabled
      Commands:
          cd
          set
          show
-> set v2c=enabled
```

6. Activate the NM2-36P Sun Datacenter InfiniBand Expansion Switch as an ASR asset in ASR Manager.

   a. Using SSH, log into the master management node.

   b. Activate the asset by adding its ILOM IP to the ASR configuration with the following command:

   ```
   # /opt/SUNWswasr/bin/asr activate_asset –i <asset_ilom_ip>
   ```

7. Repeat this procedure for the second InfiniBand switch in order to activate it as an ASR asset.

5.4.3 Activating ASR on the ZFS Storage Appliance

The ZFS Storage Appliance differs from the other ASR assets because it runs its own ASR Manager, and relays its ASR data to the Oracle backend systems through the outbound connection of the master management node. To achieve this, Oracle Private Cloud Appliance relies on the `tinyproxy` HTTP and HTTPS proxy daemon. ASR requires `tinyproxy` version 1.8.3 or later to be installed and properly configured on both management nodes.

Configuring tinyproxy for ASR

1. Using SSH and an account with superuser privileges, log into the master management node.

   Note

   The data center IP address used in this procedure is an example.

   ```
   # ssh root@10.100.1.101
   root@10.100.1.101's password:
   [root@ovcamn05r1 ~]#
   ```

2. Stop the `tinyproxy` service and verify the version installed on the system.

   ```
   # service tinyproxy stop
   Stopping tinyproxy: [ OK ]
   # rpm -q tinyproxy
   tinyproxy-1.8.2-1.el6.x86_64
   ```

3. If the installed tinyproxy is older than version 1.8.3, upgrade it to the version referenced in the support note with Doc ID 1645856.1.

   Download the tinyproxy rpm to a temporary directory on the management node and run the upgrade command.

   ```
   # cd /tmp
   # ls -al tinyproxy*
   -rw-r--r-- 1 root root 61224 Apr 4 13:44 tinyproxy-1.8.3-1.el6.x86_64.rpm
   # rpm -Uvh tinyproxy-1.8.3-1.el6.x86_64.rpm
   warning: tinyproxy-1.8.3-1.el6.x86_64.rpm: Header V3 RSA/SHA256 Signature,
   key ID ec551f03: NOKEY
   Preparing...                ########################################### [100%]
      1:tinyproxy              ########################################### [100%]
   ```

4. Make the necessary changes to the tinyproxy configuration to enable ASR traffic through the management node.

   a. Open the file `/etc/tinyproxy/tinyproxy.conf` for editing.

   b. Navigate to the section that defines on which interfaces tinyproxy listens. It starts with "`# Listen:`".

Comment out the line that binds tinyproxy to the interface in the Oracle VM management network by adding a hash character (#) in front of it.

```
# Listen: If you have multiple interfaces this allows you to bind to
# only one. If this is commented out, tinyproxy will bind to all
# interfaces present.
#
#Listen 192.168.140.4
```

c. Scroll down to the section Customization of authorization controls. It starts with "`# Allow:`".

Add two lines to allow traffic from the internal appliance management network and Oracle VM management network.

```
# Allow: Customization of authorization controls. If there are any
# access control keywords then the default action is to DENY. Otherwise,
# the default action is ALLOW.
#
# The order of the controls are important. All incoming connections are
# tested against the controls based on order.
#
#Allow 127.0.0.1
Allow 192.168.4.0/24
Allow 192.168.140.0/24
```

d. Save and close the file `/etc/tinyproxy/tinyproxy.conf`.

5. Restart the `tinyproxy` service and verify the installed version.

```
# service tinyproxy start
Starting tinyproxy: [ OK ]
# tinyproxy -version
tinyproxy 1.8.3
```

6. Repeat this procedure on the other management node.

Activating the ZFS Storage Appliance

For detailed information about enabling ASR on the ZFS Storage Appliance, refer to the support note with Doc ID 1285455.1.

You must set up the ZFS Storage Appliance to relay its ASR data through the ASR Manager on the master management node, using the following proxy settings:

- **inet addr:** 192.168.4.216
- **broadcast:** 192.168.4.255
- **mask:** 255.255.255.0
- **port:** 8888

Chapter 6 Servicing Oracle Private Cloud Appliance Components

Table of Contents

6.1 Replaceable Components 81
6.1.1 Rack Components 82
6.1.2 Oracle Server X5-2 Components 83
6.1.3 Sun Server X4-2 Components 84
6.1.4 Sun Server X3-2 Components 84
6.1.5 Oracle ZFS Storage Appliance ZS3-ES Components 85
6.1.6 Sun ZFS Storage Appliance 7320 Components 86
6.1.7 Oracle Switch ES1-24 Components 88
6.1.8 NM2-36P Sun Datacenter InfiniBand Expansion Switch Components 88
6.1.9 Oracle Fabric Interconnect F1-15 Components 89
6.2 Preparing Oracle Private Cloud Appliance for Service 89
6.3 Servicing the Oracle Private Cloud Appliance Rack System 90
6.3.1 Powering Down Oracle Private Cloud Appliance (When Required) 90
6.3.2 Service Procedures for Rack System Components 91
6.4 Servicing an Oracle Server X5-2 92
6.4.1 Powering Down Oracle Server X5-2 for Service (When Required) 92
6.4.2 Service Procedures for Oracle Server X5-2 Components 93
6.5 Servicing a Sun Server X4-2 94
6.5.1 Powering Down Sun Server X4-2 for Service (When Required) 94
6.5.2 Service Procedures for Sun Server X4-2 Components 96
6.6 Servicing a Sun Server X3-2 97
6.6.1 Powering Down Sun Server X3-2 for Service (When Required) 97
6.6.2 Service Procedures for Sun Server X3-2 Components 98
6.7 Servicing the Oracle ZFS Storage Appliance ZS3-ES 99
6.7.1 Powering Down the Oracle ZFS Storage Appliance ZS3-ES for Service (When Required) 99
6.7.2 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components 101
6.8 Servicing the Sun ZFS Storage Appliance 7320 102
6.8.1 Powering Down the Sun ZFS Storage Appliance 7320 for Service (When Required) 103
6.8.2 Service Procedures for Sun ZFS Storage Appliance 7320 Components 104
6.9 Servicing an Oracle Switch ES1-24 105
6.9.1 Powering Down the Oracle Switch ES1-24 for Service (When Required) 105
6.9.2 Service Procedures for Oracle Switch ES1-24 Components 105
6.10 Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch 106
6.10.1 Powering Down the NM2-36P Sun Datacenter InfiniBand Expansion Switch for Service (When Required) 106
6.10.2 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components 106
6.11 Servicing an Oracle Fabric Interconnect F1-15 107
6.11.1 Powering Down the Oracle Fabric Interconnect F1-15 for Service (When Required) 107
6.11.2 Service Procedures for Oracle Fabric Interconnect F1-15 Components 107

This chapter contains an overview of replaceable components in your Oracle Private Cloud Appliance (PCA), and provides servicing instructions for customer-replaceable units.

6.1 Replaceable Components

According to Oracle's Component Replacement Policy, the replaceable components in your system are designated as either field-replaceable units (FRUs) or customer-replaceable units (CRUs).

- A part designated as a FRU must be replaced by an Oracle-qualified service technician.
- A part designated as a CRU can be replaced by a person who is not an Oracle-qualified service technician.

All CRUs and FRUs are listed and identified in this chapter, but the servicing instructions included in this Oracle Private Cloud Appliance Administrator's Guide are focused primarily on CRUs. For FRU replacement, please contact Oracle.

6.1.1 Rack Components

The following table lists the replaceable components of the Oracle PCA rack.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.1 Replaceable Oracle Private Cloud Appliance Rack Components

Component Description	FRU/CRU	Hot-Swap
Sun Rack II 1242:		
Jumper Cable C13-C14, 2m	FRU	Yes
Jumper Cable C19-C20, 1m	FRU	Yes
10Gbps QSFP to QSFP Cable, 3m	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 10ft, Blue	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 1m, Grey	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 7ft, Black	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 7ft, Green	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 7ft, Yellow	FRU	Yes
Ethernet Cable, Category 6A, RJ45 to RJ45, 10ft, Black	FRU	Yes
Ethernet Cable, Category 6A, RJ45 to RJ45, 10ft, Blue	FRU	Yes
1U/2U Screw-Mount Slide Rail Kit	FRU	
1U/2U Cable Management Arm (Snap-in)	FRU	
10Gbps SFP+ TwinX Cable, 1m	FRU	Yes
1.25Gb/Sec Copper SFP Transceiver	FRU	
10Gbps QSFP to QSFP Cable, 1m	FRU	Yes
Power Distribution Units (PDUs):		
22KVA Single-Phase PDU, North America	FRU	Yes

Component Description	FRU/CRU	Hot-Swap
22KVA Single-Phase PDU, International	FRU	Yes
24KVA Three-Phase PDU, North America	FRU	Yes
24KVA Three-Phase PDU, International	FRU	Yes

For rack-level component servicing instructions, see Section 6.3, “Servicing the Oracle Private Cloud Appliance Rack System”.

6.1.2 Oracle Server X5-2 Components

The following table lists the replaceable components of the Oracle Server X5-2 management and compute nodes.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.2 Replaceable Oracle Server X5-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
Eighteen-core Intel Xeon processor E5-2699 v3 series (2.3 GHz), 145W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
32GB DDR4-2133 Load Reduced DIMM	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3)	FRU	No
8GB USB 2.0 Flash Drive	FRU	No
8-Port 12Gbps SAS-3 RAID PCI Express HBA	FRU	No
1U/2U Remote Battery Assembly	CRU	No
Cable Kit	FRU	No

For Oracle Server X5-2 component servicing instructions, see Section 6.4, “Servicing an Oracle Server X5-2”.

6.1.3 Sun Server X4-2 Components

The following table lists the replaceable components of the Sun Server X4-2 management and compute nodes.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.3 Replaceable Sun Server X4-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
2.6GHz Intel 8-core Xeon E5-2650, 95W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
16GB DDR3-1600 DIMM, 1.35V	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3)	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
8-Port 6Gbps SAS-2 RAID PCI Express HBA, B4 ASIC	FRU	No
1U/2U Remote Battery Assembly	CRU	No
Cable Kit	FRU	No

For Sun Server X4-2 component servicing instructions, see Section 6.5, “Servicing a Sun Server X4-2”.

6.1.4 Sun Server X3-2 Components

The following table lists the replaceable components of the Sun Server X3-2 management and compute nodes.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.4 Replaceable Sun Server X3-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
2.2GHz Intel 8-core Xeon E5-2660, 95W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
900GB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
16GB DDR3-1600 DIMM, 1.35V	FRU	No
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
8-Port 6Gbps SAS-2 RAID PCI Express HBA, B4 ASIC	FRU	No
1U/2U Remote Battery Assembly	CRU	No

For Sun Server X3-2 component servicing instructions, see Section 6.6, “Servicing a Sun Server X3-2”.

6.1.5 Oracle ZFS Storage Appliance ZS3-ES Components

The following table lists the replaceable components of the Oracle ZFS Storage Appliance ZS3-ES.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.5 Replaceable Oracle ZFS Storage Appliance ZS3-ES Components

Component Description	FRU/CRU	Hot-Swap
Oracle ZFS Storage Appliance ZS3-ES Storage Head:		
2.1GHz Intel 8-Core Xeon E5-2658, 95W	FRU	No
Pre-greased CPU Heatsink	FRU	No

Component Description	FRU/CRU	Hot-Swap
16GB DDR-1600 DIMM, 1.35V	FRU	No
1.6TB SAS Solid State Drive Assembly	CRU	Yes
900GB - 10000 RPM SAS Disk Assembly	CRU	Yes
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
1-Slot PCI Express Riser Assembly	FRU	No
Power Distribution Board	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
Cable Kit	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
System Board Assembly	FRU	No
1U/2U Remote Battery Assembly	CRU	No
Type A256 600 Watt AC Input Power Supply	CRU	Yes
Cluster Heartbeat Assembly	FRU	No
8-Port 6Gbps SAS-2 RAID HBA	FRU	No
8-Port 6Gbps SAS-2 PCI Express HBA (LSI)	FRU	No
Oracle Storage DE2-24P Disk Shelf:		
580 Watt AC Input Power Supply	FRU	Yes
6Gbps SAS-2 I/O Controller Module	FRU	Yes
2RU Chassis Assembly with Midplane	FRU	No
4X Mini SAS Cable, SFF-8088 to SFF-8088, 2M	FRU	
4X Mini SAS Cable, SFF-8088 to SFF-8088, 0.5M	FRU	
DE2-24P Mounting Rail Kit	FRU	
900GB 10000 RPM SAS Disk Drive Assembly	CRU	Yes
73GB SAS Solid State Drive Assembly	CRU	Yes

For Oracle ZFS Storage Appliance ZS3-ES component servicing instructions, see Section 6.7, “Servicing the Oracle ZFS Storage Appliance ZS3-ES”.

6.1.6 Sun ZFS Storage Appliance 7320 Components

The following table lists the replaceable components of the Sun ZFS Storage Appliance 7320.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.6 Replaceable Sun ZFS Storage Appliance 7320 Components

Component Description	FRU/CRU	Hot-Swap
Sun ZFS 7320 Storage Head:		
2.4GHz Intel Quad-Core Xeon E5620, 12MB, 80W	FRU	No
Xeon Heatsink	FRU	No
8GB Registered DDR3L-1333/DDR3L-1600 DIMM, 1.35V	FRU	No
512GB Solid State Drive SATA-2 Assembly	CRU	Yes
500GB - 10000 RPM SATA Disk Assembly with 1 bracket	CRU	Yes
USB Assembly	FRU	Yes
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
1-Slot x8 PCI Express Riser Assembly	FRU	No
1-Slot x16 PCI Express Riser Assembly	FRU	No
Power Distribution Board	FRU	No
8-Slot Disk Backplane, SATA DVD	FRU	No
PDB to System Board Ribbon Cable	FRU	
SFF8087 to SFF8087 Mini-SAS Cable, 690mm	FRU	
6-Pin Fan Power Cable	FRU	
Fan Data Ribbon Cable	FRU	
Bus Bar Set	FRU	
Fan Board Assembly	FRU	
Connector Board Assembly, SATA DVD	FRU	
Fan Module	CRU	Yes
System Board Assembly	FRU	No
3V Lithium Coin Cell Battery	FRU	No
Type A247A 760 Watt AC Input Power Supply	CRU	Yes
Cluster Heartbeat Assembly	FRU	
8-Port 6Gbps SAS-2 RAID HBA	FRU	No
Oracle Storage DE2-24P Disk Shelf:		
580 Watt AC Input Power Supply	CRU	Yes
6Gbps SAS-2 I/O Controller Module	FRU	Yes
2RU Chassis Assembly with Midplane	FRU	No
4X Mini SAS Cable, SFF-8088 to SFF-8088, 2M	FRU	
4X Mini SAS Cable, SFF-8088 to SFF-8088, 0.5M	FRU	
DE2-24P Mounting Rail Kit	FRU	
900GB 10000 RPM SAS Disk Drive Assembly	CRU	Yes

Component Description	FRU/CRU	Hot-Swap
73GB SAS Solid State Drive Assembly	CRU	Yes

For Sun ZFS Storage Appliance 7320 component servicing instructions, see Section 6.8, “Servicing the Sun ZFS Storage Appliance 7320”.

6.1.7 Oracle Switch ES1-24 Components

The following table lists the replaceable components of the Oracle Switch ES1-24.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.7 Replaceable Oracle Switch ES1-24 Components

Component Description	FRU/CRU	Hot-Swap
24-Port ES1-24 Switch Assembly	FRU	No
Rear-to-Front Airflow Fan Module	CRU	Yes
Type A247A 760 Watt AC Input Power Supply	CRU	Yes

For Oracle Switch ES1-24 component servicing instructions, see Section 6.9, “Servicing an Oracle Switch ES1-24”.

6.1.8 NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

The following table lists the replaceable components of the Sun Datacenter InfiniBand Expansion Switch NM2-36P.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.8 Replaceable NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

Component Description	FRU/CRU	Hot-Swap
Datacenter InfiniBand Switch 36 Subassembly	FRU	No
Type A247A 760 Watt AC Input Power Supply	CRU	Yes
Rear-to-Front Airflow Fan Module	CRU	Yes

For NM2-36P Sun Datacenter InfiniBand Expansion Switch component servicing instructions, see Section 6.10, “Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch”.

6.1.9 Oracle Fabric Interconnect F1-15 Components

The following table lists the replaceable components of the Oracle Fabric Interconnect F1-15.

Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA components list in the Oracle System Handbook.

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.9 Replaceable Oracle Fabric Interconnect F1-15 Components

Component Description	FRU/CRU	Hot-Swap
F1-15 Power Supply	FRU	Yes
QDR Fabric Board	FRU	No
2U/4U Front Panel G2 (Com-X i7)	FRU	No
F1-15 I/O Management Module	FRU	No
F1-15 Fan Tray	FRU	Yes
Quad Port 10 Gigabit Ethernet (GbE) Module	FRU	Yes
Dual Port 2 × 8 Gigabit Fibre Channel I/O Module	FRU	Yes
F1-15 Chassis without Power Supply, Fan, Fabric Board, Front Panel	FRU	No

For Oracle Fabric Interconnect F1-15 component servicing instructions, see Section 6.11, “Servicing an Oracle Fabric Interconnect F1-15”.

6.2 Preparing Oracle Private Cloud Appliance for Service

This section describes safety considerations and prerequisites for component replacement procedures.

Safety Precautions

For your protection, observe the following safety precautions when servicing your equipment:

- Follow all standard cautions, warnings, and instructions marked on the equipment and described in the following documents:
  - The printed document *Important Safety Information for Sun Hardware Systems (7063567)*
  - The *Oracle Private Cloud Appliance Safety and Compliance Guide (E56838-03)*
- Follow the safety guidelines described in the *Oracle Private Cloud Appliance Installation Guide (E56052-03)*:
  - Electrical Power Requirements
  - Rack-mount Safety Precautions

- Follow the electrostatic discharge safety practices as described in this section.
- Disconnect all power supply cords before servicing components.

Electrostatic Discharge Safety

Devices that are sensitive to electrostatic discharge (ESD), such as motherboards, PCIe cards, drives, processors, and memory cards require special handling.

Caution

Equipment Damage

Take antistatic measures and do not touch components along their connector edges.

- **Use an antistatic wrist strap.**

  Wear an antistatic wrist strap and use an antistatic mat when handling components such as drive assemblies, boards, or cards. When servicing or removing rack node components, attach an antistatic strap to your wrist and then to a metal area on the chassis. Then disconnect the power cords from the component. Following this practice equalizes the electrical potentials between you and the component.

  An antistatic wrist strap is *not* included in the Oracle PCA shipment.

- **Use an antistatic mat.**

  Place ESD-sensitive components such as the motherboard, memory, and other PCB cards on an antistatic mat.

  The following items can be used as an antistatic mat:

  - Antistatic bag used to wrap an Oracle replacement part
  - An ESD mat (orderable from Oracle)
  - A disposable ESD mat (shipped with some replacement parts or optional system components)

6.3 Servicing the Oracle Private Cloud Appliance Rack System

This section provides instructions to service replaceable components (CRUs/FRUs) in the appliance rack. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.3.1 Powering Down Oracle Private Cloud Appliance (When Required)

Some service procedures may require you to power down the Oracle PCA. Perform the following steps to manually power down the system.

Powering down the system

1. Press the Power button on the compute nodes in order to shut them down gracefully.
2. Press the Power button on the management nodes in order to shut them down gracefully.

   Once the servers are powered off, you can proceed to power off the storage appliance.
3. Press the Power button on the storage server heads attached to the chassis of the storage device.

4. Press the Power button on both Fabric Interconnects.
5. Toggle the rack Power switches to the Off position.

Note

The Ethernet and InfiniBand switches do not have power switches. They power off when power is removed, by way of the power distribution unit (PDU) or at the breaker in the data center.

Returning the system to operation

1. Toggle the power distribution unit (PDU) circuit breakers of both PDUs to the On position.
2. Wait at least two minutes to allow the PDUs to complete their power-on sequence.

   The Ethernet and InfiniBand switches are powered on with the PDUs.
3. Press the Power button on each Fabric Interconnect.

   Wait until the Status LED stops blinking and is solid green, indicating the system has finished booting.
4. Press the Power button on the storage server heads.

   Wait approximately two minutes until the power-on self-test completes, and the Power/OK LED on the front panel lights and remains lit.
5. Press the Power button on the management nodes.

   The management node that completes booting first assumes the master role.
6. When the management nodes are up, press the Power button on the compute nodes.

   When all compute nodes are up, verify the status of all system components in Oracle VM Manager.

   If no components are in error state, the appliance is ready to resume normal operation.

6.3.2 Service Procedures for Rack System Components

For parts that are not hot-swappable, power down the Oracle PCA before starting the service procedure. Generally speaking, hot-swappable components can be serviced without specific additional steps.

Table 6.10 Service Instructions for Rack System Components

Replaceable Part(s)	Hot-Swap	Instructions
Power cables		
Ethernet cables		
InfiniBand cables		
SFP transceivers (Oracle-qualified service technician only)		
Cable management arms (CMAs) (Oracle-qualified service technician only)		For removal and installation of a cable management arm, refer to the Oracle Server X5-2 Installation Guide (part no. E48314). • "Remove the Cable Management Arm"

Replaceable Part(s)	Hot-Swap	Instructions
		• "Install the Cable Management Arm"
Slide rails (Oracle-qualified service technician only)		To service the slide rails, the server must be removed from the rack. For instructions, refer to the Oracle Server X5-2 Service Manual (part no. E48320). • "Remove the Server From the Rack" • "Reinstall the Server Into the Rack" For slide rail installation instructions, refer to the section Attach Tool-less Slide-Rail Assemblies in the Oracle Server X5-2 Installation Guide (part no. E48314). To remove the slide rails, reverse the installation steps.

6.4 Servicing an Oracle Server X5-2

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Server X5-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in Section 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.4.1 Powering Down Oracle Server X5-2 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Server X5-2 to be powered down, follow these instructions:

Note

The management nodes are not placed in maintenance node for servicing. If you need to power down the master management node, bring it offline as described below and wait for the other management node to take over the master role. If you need to power down the secondary management node, no additional steps are required.

Placing a compute node into maintenance mode

Before an Oracle Server X5-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. For details, refer to the section "Placing an Oracle VM Server into Maintenance Mode" in the Oracle VM User's Guide.

1. Log in to the Oracle VM Manager Web UI.

   For details, refer to the section "Section 4.2, "Logging in to the Oracle VM Manager Web UI"" in the Oracle Private Cloud Appliance Administrator's Guide.

   a. Enter the following address in a Web browser: https://*manager-vip*:7002/ovm/console.

      Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

   b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

   The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Oracle Server X5-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter "Preparing for Service" in the Oracle Server X5-2 Service Manual (part no. E48320).

1. Power down the server gracefully whenever possible.

   The easiest way is to press and quickly release the Power button on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.

Caution

System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Disconnect the power cables and data cables from the server.
4. Extend the server to the maintenance position.
5. Most service operations can be performed while the server is in the maintenance position.

   However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.

Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter "Returning the Server to Operation" in the Oracle Server X5-2 Service Manual (part no. E48320).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.4.2 Service Procedures for Oracle Server X5-2 Components

For parts that are not hot-swappable, power down the Oracle Server X5-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This

protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.11 Service Procedures for Oracle Server X5-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037394.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037409.html#scrolltoc
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40003f01423753.html#scrolltoc

6.5 Servicing a Sun Server X4-2

This section provides instructions to service replaceable components (CRUs/FRUs) in a Sun Server X4-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.5.1 Powering Down Sun Server X4-2 for Service (When Required)

If you need to execute a service procedure that requires the Sun Server X4-2 to be powered down, follow these instructions:

Note

The management nodes are not placed in maintenance node for servicing. If you need to power down the master management node, bring it offline as described below and wait for the other management node to take over the master role. If you

need to power down the secondary management node, no additional steps are required.

Placing a compute node into maintenance mode

Before a Sun Server X4-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. For details, refer to the section "Placing an Oracle VM Server into Maintenance Mode" in the Oracle VM User's Guide.

1. Log in to the Oracle VM Manager Web UI.

   For details, refer to the section "Section 4.2, "Logging in to the Oracle VM Manager Web UI"" in the Oracle Private Cloud Appliance Administrator's Guide.

   a. Enter the following address in a Web browser: https://*manager-vip*:7002/ovm/console.

      Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

   b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

   The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

   The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Sun Server X4-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter "Preparing for Service" in the Sun Server X4-2 Service Manual (part no. E38041).

1. Power down the server gracefully whenever possible.

   The easiest way is to press and quickly release the Power button on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.

   Caution

   System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Disconnect the power cables and data cables from the server.

4. Extend the server to the maintenance position.

5. Most service operations can be performed while the server is in the maintenance position.

   However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.

Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter "Returning the Server to Operation" in the Sun Server X4-2 Service Manual (part no. E38041).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.5.2 Service Procedures for Sun Server X4-2 Components

For parts that are not hot-swappable, power down the Sun Server X4-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.12 Service Procedures for Sun Server X4-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000f91037394.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000f91037409.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40003f01423753.html#scrolltoc

6.6 Servicing a Sun Server X3-2

This section provides instructions to service replaceable components (CRUs/FRUs) in a Sun Server X3-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.6.1 Powering Down Sun Server X3-2 for Service (When Required)

If you need to execute a service procedure that requires the Sun Server X3-2 to be powered down, follow these instructions:

Note

The management nodes are not placed in maintenance node for servicing. If you need to power down the master management node, bring it offline as described below and wait for the other management node to take over the master role. If you need to power down the secondary management node, no additional steps are required.

Placing a compute node into maintenance mode

Before a Sun Server X3-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. For details, refer to the section “Placing an Oracle VM Server into Maintenance Mode” in the Oracle VM User's Guide.

1. Log in to the Oracle VM Manager Web UI.

   For details, refer to the section “Section 4.2, “Logging in to the Oracle VM Manager Web UI”” in the Oracle Private Cloud Appliance Administrator's Guide.

   a. Enter the following address in a Web browser: https://*manager-vip*:7002/ovm/console.

      Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

   b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

   The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

   The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Sun Server X3-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter "Preparing for Service" in the Sun Server X3-2 Service Manual (part no. E22313).

1. Power down the server gracefully whenever possible.

   The easiest way is to press and quickly release the Power button on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.

Caution

System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Extend the server to the maintenance position.
4. Disconnect the power cables and data cables from the server.
5. Most service operations can be performed while the server is in the maintenance position.

   However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.

Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter "Returning the Server to Operation" in the Sun Server X3-2 Service Manual (part no. E22313).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Reconnect data cables and power cords.
4. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
5. Power on the server.

6.6.2 Service Procedures for Sun Server X3-2 Components

For parts that are not hot-swappable, power down the Sun Server X3-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.13 Service Procedures for Sun Server X3-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000f91037394.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000f91037409.html#scrolltoc
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40003f01423753.html#scrolltoc

6.7 Servicing the Oracle ZFS Storage Appliance ZS3-ES

This section provides instructions to service replaceable components (CRUs/FRUs) in the Oracle ZFS Storage Appliance ZS3-ES. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.7.1 Powering Down the Oracle ZFS Storage Appliance ZS3-ES for Service (When Required)

If you need to execute a service procedure that requires the Oracle ZFS Storage Appliance ZS3-ES to be powered down, follow these instructions:

Powering down the storage head/controller

Performing a graceful shutdown ensures that data is saved and not corrupted, and that resources are assigned to the other controller in the storage head cluster. This is the preferred method for powering down a controller for component replacement.

1. Ensure that Ethernet cables are connected from your network to the `NET-0` port on the back of each server.
2. Direct your web browser to the server to be serviced by using either the IP address or host name assigned to the NET-0 port as follows: `https://ipaddress:215`.
3. Log in as root, using the system-wide Oracle PCA password.

4. Go to **Maintenance**, then select **Hardware**.
5. Click the **Show Details** link for the server.
6. Click the Power icon for the server and select **Power off** from the pull-down list.

If graceful shutdown is not possible, use the power button.

Caution

This task forces the main power off. You might corrupt or lose system data, or lose the server configuration (the resources assigned to it) during an immediate power down.

1. Press and quickly release the Power button on the front panel.

   This action causes an orderly shutdown of the operating system, and the server enters the standby power mode.

2. If the server did not respond or you need a more immediate shutdown, press and hold the Power button for four seconds.

   This forces the main power off and enters the standby power mode immediately. When the main power is off, the Power/OK LED on the front panel begins flashing, indicating that the server is in standby power mode.

If neither graceful shutdown nor emergency shutdown using the power button is possible, for example because you are not physically located at the system, use the ILOM to perform an emergency shutdown. Choose one of the following options:

Caution

This task forces the main power off. You might corrupt or lose system data, or lose the server configuration (the resources assigned to it) during an immediate power down.

- Log in to the **Oracle ILOM web interface**.

  In the left pane, click Host Management > Power Control, and in the Actions list click **Immediate Power Off**.

  Click Save, and then click OK.

- Log in to the **Oracle ILOM command-line interface** (CLI).

  At the CLI prompt, type the following command: `stop -f /System`.

Powering down the disk shelf

Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place. Powering down or removing all SAS chains from a disk shelf will cause the controllers to panic to prevent data loss. To avoid this, shut down the controllers before decommissioning the shelf.

1. Stop all input and output to and from the disk shelf.
2. Wait approximately two minutes until all disk activity indicators have stopped flashing.
3. Place the power supply on/off switches to the "O" off position.

4. Disconnect the power cords from the external power source.

Powering on the storage appliance

The disk shelf must not be operated without all components in place.

1. Reconnect the disk shelf power and data cables you removed to service a component.
2. Place the power supply on/off switches on the disk shelf to the "I" on position.
3. Wait several minutes until the boot process is complete, at which time the Power LED should be solid green.
4. Connect the storage head power and data cables you removed to service a component.
5. Power on the server by pressing the Power button on the front panel.

   If you are not physically located at the system, use either of these ILOM methods instead:

   - Log in to the **Oracle ILOM web interface**.

     In the left pane, click Host Management > Power Control, and in the Actions list click **Power On**.

   - Log in to the **Oracle ILOM command-line interface** (CLI).

     At the CLI prompt, type the following command: `start /System`.
6. Wait approximately two minutes until the power-on self-test (POST) code checkpoint tests have completed, and the Power/OK LED on the front panel lights and remains lit.
7. If you performed a graceful shutdown earlier, return resources to the server that was just serviced.

   a. Log into the web UI for the server that was not serviced.

   b. Go to Configuration > Cluster.

   c. Click **Failback**.

Note

For information about configuring the clustered servers and attached disk shelves, see the “Oracle ZFS Storage System Administration Guide” for the appropriate software release.

6.7.2 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components

For parts that are not hot-swappable, power down the Oracle ZFS Storage Appliance ZS3-ES before starting the service procedure.

Warning

If you need to execute a service procedure that interrupts the connection between virtual machines and their virtual disks, shut down the virtual machines in Oracle VM Manager prior to servicing the storage hardware. Disconnecting a running virtual machine from its disks may cause data corruption.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to

each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.14 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components

Replaceable Part(s)	Hot-Swap	URL
Storage head hard drives	Yes	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000091011460.html#scrolltoc
Disk shelf drives	Yes	Refer to the section “Replacing a Drive” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Fan modules	Yes	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000091014194.html#scrolltoc
Storage head power supplies	Yes	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000091014153.html#scrolltoc
Disk shelf power supplies	Yes	Refer to the section “Replacing a Power Supply” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Memory modules (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000f91037394.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000f91037409.html#scrolltoc
Internal USB flash drive (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40003f01423753.html#scrolltoc
Disk shelf I/O modules (Oracle-qualified service technician only)	Yes	Refer to the section “Replacing an I/O Module” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Disk shelf SIM boards (Oracle-qualified service technician only)	Yes	Refer to the section “Replacing a SIM Board” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc

6.8 Servicing the Sun ZFS Storage Appliance 7320

This section provides instructions to service replaceable components (CRUs/FRUs) in the Sun ZFS Storage Appliance 7320. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.8.1 Powering Down the Sun ZFS Storage Appliance 7320 for Service (When Required)

If you need to execute a service procedure that requires the Sun ZFS Storage Appliance 7320 to be powered down, follow these instructions:

Powering down the storage head/controller

Powering down or removing all SAS chains from a disk shelf will cause the controllers to panic to prevent data loss. To avoid this, shut down the controllers before decommissioning the shelf.

1. Log in to the BUI.
2. Click the Power icon on the left side of the masthead.

If the BUI is not accessible, select one of the following options:

Note

In a configuration with clustered storage heads, always shut down the standby head before the active head.

- SSH into the appliance and issue the `maintenance system poweroff` command.
- SSH or serial console into the service processor (SP) and issue the `stop /SYS` command.
- Use a pen or non-conducting pointed object to press and release the Power button on the front panel.

Caution

To initiate emergency shutdown during which all applications and files will be closed abruptly without saving, press and hold the power button for at least four seconds until the Power/OK status indicator on the front panel flashes, indicating that the storage controller is in standby power mode.

Powering down the disk shelf

Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place. Powering down or removing all SAS chains from a disk shelf will cause the controllers to panic to prevent data loss. To avoid this, shut down the controllers before decommissioning the shelf.

1. Stop all input and output to and from the disk shelf.
2. Wait approximately two minutes until all disk activity indicators have stopped flashing.
3. Place the power supply on/off switches to the "O" off position.
4. Disconnect the power cords from the external power source.

Powering on the storage appliance

The disk shelf must not be operated without all components in place.

1. Reconnect the disk shelf power and data cables you removed to service a component.
2. Place the power supply on/off switches on the disk shelf to the "I" on position.

3. Wait several minutes until the boot process is complete, at which time the Power LED should be solid green.

4. Connect the storage head power cables and wait approximately two minutes until the Power/OK LED on the front panel next to the Power button lights and remains lit.

6.8.2 Service Procedures for Sun ZFS Storage Appliance 7320 Components

For parts that are not hot-swappable, power down the Sun ZFS Storage Appliance 7320 before starting the service procedure.

Warning

If you need to execute a service procedure that interrupts the connection between virtual machines and their virtual disks, shut down the virtual machines in Oracle VM Manager prior to servicing the storage hardware. Disconnecting a running virtual machine from its disks may cause data corruption.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.15 Service Procedures for Sun ZFS Storage Appliance 7320 Components

Replaceable Part(s)	Hot-Swap	URL
Storage head HDDs or SSDs	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardwar
Disk shelf drives	Yes	Refer to the section “Replacing a Drive” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Fan modules	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardwar
Storage head power supplies	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardwar
Disk shelf power supplies	Yes	Refer to the section “Replacing a Power Supply” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Memory modules (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardwar
PCI Express risers and cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardwar
Battery	No	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardwar
System indicator boards (Oracle-qualified service technician only)	Yes	http://docs.oracle.com/cd/E26765_01/html/E26399/maintenance__hardware__details__7x20.html#maintenance__hardwar

Replaceable Part(s)	Hot-Swap	URL
Disk shelf I/O modules (Oracle-qualified service technician only)	Yes	Refer to the section “Replacing an I/O Module” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Disk shelf SIM boards (Oracle-qualified service technician only)	Yes	Refer to the section “Replacing a SIM Board” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc

6.9 Servicing an Oracle Switch ES1-24

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Switch ES1-24. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.9.1 Powering Down the Oracle Switch ES1-24 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Switch ES1-24 to be powered down, follow these instructions:

Powering down the switch

1. To power down an individual power supply, remove its power cord.
2. To power down the switch, remove the power cords from both power supplies.

Returning the switch to operation

1. Reconnect the power cords to both power supplies.
2. Verify that the switch has power by checking the status LEDs.

   The AC LED lights green to indicate the power supply is connected to line power. A moment later, the OK LED lights green to indicate the power supply is fully operational.

6.9.2 Service Procedures for Oracle Switch ES1-24 Components

For parts that are not hot-swappable, power down the Oracle Switch ES1-24 before starting the service procedure.

Warning

Internal Ethernet connectivity is affected while the component is out of service. Please take the necessary precautions.

Caution

When replacing the entire switch assembly, begin by saving the configuration from the existing component, so that you can restore the configuration after replacement.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.16 Service Procedures for Oracle Switch ES1-24 Components

Replaceable Part(s)	Hot-Swap	URL
Power supplies	Yes	http://docs.oracle.com/cd/E39109_01/html/E39116/z40000349112.html#scrolltoc
Fan module	Yes	http://docs.oracle.com/cd/E39109_01/html/E39116/z40000369112.html#scrolltoc

6.10 Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch

This section provides instructions to service replaceable components (CRUs/FRUs) in a NM2-36P Sun Datacenter InfiniBand Expansion Switch. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.10.1 Powering Down the NM2-36P Sun Datacenter InfiniBand Expansion Switch for Service (When Required)

If you need to execute a service procedure that requires the NM2-36P Sun Datacenter InfiniBand Expansion Switch to be powered down, follow these instructions:

Powering down the switch

1. To power down an individual power supply, remove its power cord.
2. To power down the switch, remove the power cords from both power supplies.

Returning the switch to operation

1. Reconnect the power cords to both power supplies.
2. Verify that the switch has power by checking the status LEDs.

   The AC LED lights green to indicate the power supply is connected to line power. A moment later, the OK LED lights green to indicate the power supply is fully operational.

6.10.2 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

For parts that are not hot-swappable, power down the NM2-36P Sun Datacenter InfiniBand Expansion Switch before starting the service procedure.

Caution

InfiniBand connectivity may be affected while the component is out of service. Please take the necessary precautions.

Caution

When replacing the entire switch assembly, begin by saving the configuration from the existing component, so that you can restore the configuration after replacement.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.17 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

Replaceable Part(s)	Hot-Swap	URL
Power supplies	Yes	http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f49112.html#scrolltoc
Fans	Yes	http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f59112.html#scrolltoc
Data cables	Yes	http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f69112.html#scrolltoc

6.11 Servicing an Oracle Fabric Interconnect F1-15

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Fabric Interconnect F1-15. Before starting any service procedure, read and follow the guidelines in Section 6.2, “Preparing Oracle Private Cloud Appliance for Service”.

6.11.1 Powering Down the Oracle Fabric Interconnect F1-15 for Service (When Required)

If you need to execute a service procedure that requires the Fabric Interconnect to be powered down, follow these instructions:

Powering down the Oracle Fabric Interconnect F1-15

1. Press the Power button to power down the Fabric Interconnect gracefully.
2. Wait for the Status LED to switch off, indicating that the component has been powered down successfully.

Returning the Oracle Fabric Interconnect F1-15 to operation

1. Press the Power button to power on the Fabric Interconnect.

   The Status LED blinks green, indicating that the system control processor is booting.
2. Wait until the Status LED is solid green.

   This indicates that the system control processor has finished booting and the Fabric Interconnect is ready for operation.

6.11.2 Service Procedures for Oracle Fabric Interconnect F1-15 Components

For parts that are not hot-swappable, power down the Oracle Fabric Interconnect F1-15 before starting the service procedure.

Caution

Management, storage, VM and external network connectivity may be affected while the Fabric Interconnect or an I/O module is out of service. Please take the necessary precautions.

Caution

When replacing the entire switch assembly, begin by saving the configuration from the existing component, so that you can restore the configuration after replacement.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle PCA. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.18 Service Procedures for Oracle Fabric Interconnect F1-15 Components

Replaceable Part(s)	Hot-Swap	URL
Power supplies (Oracle-qualified service technician only)	Yes	In the Oracle Fabric Director Hardware and Host Drivers Installation Guide, refer to the following section: Chapter 7 (page 179): “Performing Service Tasks”, page 189: “Removing VP780 Power Supplies”. Installation is the reverse of this procedure.
Fan modules (Oracle-qualified service technician only)	Yes	In the Oracle Fabric Director Hardware and Host Drivers Installation Guide, refer to the following section: Chapter 7 (page 179): “Performing Service Tasks”, page 184: “Removing VP780 Fan Modules”. Installation is the reverse of this procedure.
Fabric board (Oracle-qualified service technician only)	No	In the Oracle Fabric Director Hardware and Host Drivers Installation Guide, refer to the following section: Chapter 7 (page 179): “Performing Service Tasks”, page 202: “Replacing the VP780 QDR Fabric Board”.
Management module (Oracle-qualified service technician only)	No	In the Oracle Fabric Director Hardware and Host Drivers Installation Guide, refer to the following section: Chapter 7 (page 179): “Performing Service Tasks”, page 235: “Removing and Installing a VP780 Management Module”.
I/O modules (Oracle-qualified service technician only)	Yes	In the Oracle Fabric Director Hardware and Host Drivers Installation Guide, refer to the following section: Chapter 7 (page 179): “Performing Service Tasks”, page 224: “Removing and Installing VP780 I/O Modules”. **Caution** The latest generation I/O modules have no Eject button and are powered down from the CLI. For detailed instructions, refer to the support note with Doc ID 1518778.1.
Front panel assembly, including system control processor (Oracle-qualified service technician only)	No	In the Oracle Fabric Director Hardware and Host Drivers Installation Guide, refer to the following section: Chapter 7 (page 179): “Performing Service Tasks”, page 250: “Removing and Installing a VP780 Front Panel Assembly”.

Replaceable Part(s)	Hot-Swap	URL	
			Caution When replacing the system control processor (SCP), which is included in the front panel assembly, begin by saving the configuration from the existing SCP, so that you can restore the configuration after replacement.

Chapter 7 Troubleshooting

Table of Contents

7.1 Adding Proxy Settings for Oracle Private Cloud Appliance Updates 111
7.2 Setting the Oracle Private Cloud Appliance Logging Parameters 112
7.3 Configuring Data Center Switches for VLAN Traffic 112
7.4 Enabling Support Tools on Systems Provisioned with Software Release 1.0.2 113
7.5 Verifying and Re-applying Oracle VM Manager Tuning after Software Update 114
7.6 Upgrading Existing Compute Node Configuration from Release 1.0.2 115
7.7 Enabling Compute Node IPMI in Oracle VM Manager 118
7.8 Enabling SNMP Server Monitoring 118
7.9 Changing Oracle WebLogic Server Passwords 120
7.10 Replacing Default Passwords Manually 121
7.11 Using a Custom CA Certificate for SSL Encryption 123
7.12 A Compute Node Fails to Complete Provisioning 123
7.13 Oracle VM Server Pool Is Offline After Network Services Restart 125

This chapter describes how to resolve a number of common problem scenarios.

7.1 Adding Proxy Settings for Oracle Private Cloud Appliance Updates

If your data center does not provide unlimited internet access and has a proxy server in place to control HTTP, HTTPS or FTP traffic, you may need to configure your management nodes to be able to access external resources for the purpose of performing software updates.

The following instructions must be followed for each of the two management nodes in your environment.

Adding Proxy Settings for a Management Node

1. Gain command line access to the management node. Usually this is achieved using SSH and logging in as the root user with the global Oracle PCA password.

2. Use the CLI, as described in Chapter 3, *The Oracle Private Cloud Appliance Command Line Interface (CLI)*, to view or modify your proxy settings. The CLI safely reads and edits the `/etc/ovca.conf` file, to prevent the possibility of configuration file corruption.

   - To view the current values for the configurable settings in the configuration file run the CLI as follows:

     ```
     # pca-admin show system-properties
     ```

   - To set an HTTP proxy:

     ```
     # pca-admin set system-property http_proxy http://IP:PORT
     ```

     Where *IP* is the IP address of your proxy server, and *PORT* is the TCP port on which it is listening. If your proxy server expects a user name and password, this can be specified in the following format:

     ```
     http://username:password@IP:PORT
     ```

   - To set an HTTPS proxy:

     ```
     # pca-admin set system-property https_proxy https://IP:PORT
     ```

- To set an FTP proxy:

```
# pca-admin set system-property ftp_proxy ftp://IP:PORT
```

3. Setting any single parameter automatically rewrites the configuration file and the proxy settings become active immediately.

7.2 Setting the Oracle Private Cloud Appliance Logging Parameters

When troubleshooting or if you have a support query open, you may be required to change the logging parameters for your Oracle PCA. The settings for this are contained in `/etc/ovca.conf`, and can be changed using the CLI.

The following instructions must be followed for each of the two management nodes in your environment.

Changing the Oracle Private Cloud Appliance Logging Parameters for a Management Node

1. Gain command line access to the management node. Usually this is achieved using SSH and logging in as the root user with the global Oracle PCA password.

2. Use the CLI, as described in Chapter 3, *The Oracle Private Cloud Appliance Command Line Interface (CLI)*, to view or modify your appliance log settings. The CLI safely reads and edits the `/etc/ovca.conf` file, to prevent the possibility of configuration file corruption.

   - To view the current values for the configurable settings in the configuration file run the CLI as follows:

     ```
     # pca-admin show system-properties
     ```

   - To change the log level:

     ```
     # pca-admin set system-property log_level LEVEL
     ```

     Where *`LEVEL`* is one of the following: `DEBUG`, `INFO`, `WARNING`, `ERROR`, `CRITICAL`.

   - To change the log file size:

     ```
     # pca-admin set system-property log_size SIZE
     ```

     Where *`SIZE`*, expressed in MB, is a number from 1 to 512.

   - To change the number of backup log files stored:

     ```
     # pca-admin set system-property log_count COUNT
     ```

     Where *`COUNT`* is a number of files ranging from 0 to 100.

3. The new log level setting only takes effect after a management node has been rebooted or the service has been restarted by running the `service ovca restart` command on the active management node shell.

7.3 Configuring Data Center Switches for VLAN Traffic

The Oracle PCA network infrastructure supports the use of VLANs by default. For this purpose, the Oracle Fabric Interconnect F1-15s are set to trunking mode to allow tagged data traffic. You may implement VLANs for logical separation of different network segments, or to define security boundaries between networks with different applications – just as you would with physical servers instead of virtual machines.

However, to allow virtual machines hosted by the Oracle PCA to communicate with systems external to the appliance, you must update the configuration of your next-level data center switches accordingly.

- The switch ports on the receiving end of the outbound appliance connections must be part of each VLAN used within the Oracle PCA environment.
- The same ports must also be part of the network(s) connecting the external systems that your virtual machines need to access. For example, WAN connectivity implies that virtual machines are able to reach the public gateway in your data center. As an alternative to VLAN tagging, Layer 3 routing can be used to connect to the Oracle PCA.

7.4 Enabling Support Tools on Systems Provisioned with Software Release 1.0.2

Oracle PCA controller software Release 1.1.1 and later installs the support tools on the management nodes and compute nodes. However, compute nodes provisioned and discovered in an earlier release do not have the packages for these support tools. After you update the appliance software to the new release, any compute node you provision or reprovision will be installed with the Oracle PCA support tools. On running compute nodes that were provisioned with controller software Release 1.0.2, and already host virtual machines, you can install the support tools manually by following the instructions in this section.

Installing Support Tools Manually on a Compute Node

1. Make sure that the appliance software has been updated successfully to Release 2.0.x.
2. Using SSH and an account with superuser privileges, log into the compute node where you want to install the support tools.

   You can connect to a compute node from a workstation connected to the appliance management network. Alternatively, log into a management node first, and then connect to the compute node through either the appliance management network (`192.168.4.x`) or the Oracle VM management network (`192.168.140.x`). The IP address of a discovered compute node can be found in the *Servers and VMs* tab of Oracle VM Manager.

   Note

   The management node data center IP address used in this procedure is an example.

   ```
   # ssh root@10.100.1.101
   root@10.100.1.101's password:
   [root@ovcamn05r1 ~]# ssh root@192.168.4.6
   root@192.168.4.6's password:
   [root@ovcacn09r1 ~]#
   ```

3. On the compute node, mount the shared partition on the storage appliance where the software install media is unpacked.

   ```
   [root@ovcacn09r1 ~]# mkdir /tmp/shared_storage
   [root@ovcacn09r1 ~]# mount 192.168.40.1:/export/MGMT_ROOT /tmp/shared_storage
   ```

4. Locate and install the `ovca-support` package.

   ```
   [root@ovcacn09r1 ~]# ls -l /tmp/shared_storage/ovm_image/Server/ovca-support*
   -r--r--r--+ 1 root root 162726 Mar 23 16:14 /tmp/shared_storage/ovm_image/Server/ \
   ovca-support-1.1.1-106.el5.noarch.rpm
   [root@ovcacn09r1 ~]# rpm -ivh /tmp/shared_storage/ovm_image/Server/ovca-support*
   warning: /tmp/shared_storage/ovm_image/Server/ovca-support-1.1.1-106.el5.noarch.rpm:
       Header V3 DSA signature: NOKEY, key ID 1e5e0159
   ```

```
    Preparing...                ########################################### [100%]
       1:ovca-support           ########################################### [100%]
[root@ovcacn09r1 ~]#
```

5. Verify that the package has been installed correctly, for example by checking the software version as follows:

```
[root@ovcacn09r1 ~]# pca-diag version
2.0.5_54_2015-06-11
[root@ovcacn09r1 ~]#
```

6. Unmount the shared partition and remove the temporary directory you created.

```
[root@ovcacn09r1 ~]# umount /tmp/shared_storage
[root@ovcacn09r1 ~]# rmdir /tmp/shared_storage
[root@ovcacn09r1 ~]#
```

7. Log out of the compute node.

8. Repeat this procedure on all other compute nodes installed in your Oracle PCA environment.

7.5 Verifying and Re-applying Oracle VM Manager Tuning after Software Update

During a software update from Release 1.0.2 to Release 2.0.1 it may occur that certain Oracle VM Manager tuning settings are not applied properly and default settings are used instead. After updating the Oracle PCA software stack, you must verify these tuning settings, and re-apply them if necessary. Run the following procedure:

Verifying Oracle VM Manager Tuning Settings

1. Using SSH and an account with superuser privileges, log into the master management node.

Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Verify that you are logged in to the **master** management node.

```
[root@ovcamn05r1 ~]# pca-check-master
NODE: 192.168.4.3  MASTER: True
```

If the command returns `MASTER: False`, log in to the other management node and run the same command.

3. Log in to the Oracle VM shell as the `admin` user.

```
# /usr/bin/ovm_shell.sh -u admin
Password:
OVM Shell: 3.2.<version_id> Interactive Mode
---
>>>
```

4. At the Oracle VM shell prompt, enter the following command:

```
>>> OvmClient.getOvmManager().getFoundryContext().getModelManager().getMaxCacheSize()
150000
```

```
>>>
```

To exit Oracle VM shell, press `Ctrl+D`.

If the value returned is **not** `300000`, proceed with the next step.

5. From the Oracle Linux command line on the master management node, apply the required Oracle VM Manager tuning settings by running the following Oracle VM shell script as the `admin` user:

```
# /usr/bin/ovm_shell.sh -u admin -i /var/lib/ovca/ovm_scripts/ovmm_tuning.py
Password:
live events max age: 24 hours
archive events max age: 72 hours
max cache size: 150000 objects
live jobs max age: 168 hours
archive jobs max age: 14 hours
live jobs max age (after): 24 hours
archive jobs max age (after): 168 hours
live events max age (after): 3 hours
archive events max age (after): 6 hours
max cache size (after): 300000 objects
```

6. When the tuning script completes successfully, log out of the master management node.

7.6 Upgrading Existing Compute Node Configuration from Release 1.0.2

When you run a Software Update on an existing Oracle PCA installation with controller software Release 1.0.2, you must execute an additional upgrade procedure on provisioned compute nodes. The procedure is designed to reconfigure the compute node storage connections and the Oracle VM Agent, allowing the compute node to continue to function reliably as a member of the clustered server pool after the software update. By following the instructions in this section, you also make sure that previously deployed virtual machines remain fully functional when the appliance update to Release 2.0.1 or newer is complete.

Warning

First update any Release 1.0.2 system to appliance software Release 1.1.3. Execute the procedure in this section on each compute node provisioned with software Release 1.0.2, *after* the Release 1.1.3 software update on the management nodes has completed successfully. Proceed with the update to Release 2.0.1 or newer only when these steps have been completed.

Note

A large part of this procedure must be executed from Oracle VM Manager. An overview of its usage in the context of Oracle PCA, and links to detailed instructions in the Oracle VM documentation, are provided in this Administrator's Guide, in Chapter 4, *Managing the Oracle VM Virtual Infrastructure*.

Upgrading the Compute Node Configuration from Release 1.0.2 to Release 1.1.3

1. Make sure that the appliance software has been updated successfully to Release 1.1.3.

   You can verify this by logging into the master management node and entering the following command:

```
# pca-diag version
1.1.3_<build>_<date>
```

2. Log in to Oracle VM Manager.

For details, see Section 4.2, “Logging in to the Oracle VM Manager Web UI”.

3. Identify the master server in the server pool.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Info** perspective from the list in the toolbar.

      A configuration overview of the server pool is displayed.

   c. Locate the *Master Server* setting and write down which server is the pool master.

Warning

The server pool master server must be the last Oracle VM Server to go through this upgrade procedure.

4. Migrate all running virtual machines away from the compute node you want to upgrade.

   Information on migrating virtual machines is provided in the Oracle VM User's Guide section entitled Migrating Virtual Machines.

5. Place the compute node in maintenance mode.

   Information on maintenance mode is provided in the Oracle VM User's Guide section entitled Placing an Oracle VM Server into Maintenance Mode.

   a. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

      The **Edit Server** dialog box is displayed.

   b. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

      The Oracle VM Server is in maintenance mode and ready for servicing.

6. Install the `ovca-support` package by executing the steps described in Section 7.4, “Enabling Support Tools on Systems Provisioned with Software Release 1.0.2”. However, do not log out of the compute node yet.

7. From the Oracle Linux command line on the compute node, run this script:

```
# /usr/sbin/cn_upg_111.sh
```

8. When the upgrade script completes successfully, log out of the compute node and return to the Oracle VM Manager UI.

9. Run the Oracle VM Server update on the compute node in question.

   Information on the Oracle VM Server update functionality is provided in the Oracle VM User's Guide section entitled Updating and Upgrading Oracle VM Servers.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Servers** perspective from the list in the toolbar.

      A table is displayed with all servers in the server pool.

c. Select the appropriate compute node in the management pane table and click **Update Server** in the management pane toolbar.

   As part of the update procedure, the Oracle VM Server is restarted but remains in maintenance mode.

Warning

If the compute node does not reboot during the update, you must restart it from within Oracle VM Manager.

10. Take the compute node out of maintenance mode.

    a. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

       The **Edit Server** dialog box is displayed.

    b. Clear the **Maintenance Mode** check box. Click OK.

       The Oracle VM Server rejoins the server pool as a fully functioning member.

11. If the compute node is in *WARNING* status, unpresent the *Rack1-Repository* from this node, and remove the node from the *Rack1_ServerPool*.

    a. In the **Repositories** tab, select the *Rack1-Repository* in the management pane. Click **Present/ Unpresent Repository** in the management pane toolbar.

       The **Present Repository to Servers** dialog box is displayed.

    b. Select the **Servers** option at the top of the dialog box. Move the server in question out of the *Present to Server* pane. Click OK.

       The repository is unpresented from the selected Oracle VM Server.

    c. In the **Servers and VMs** tab, select *Rack1_ServerPool* and click **Edit Server Pool**.

       The **Edit Server Pool** dialog box is displayed.

    d. In the **Servers** tab of the **Edit Server Pool** dialog box, move the server in question out of the *Selected Servers* pane. Click OK.

       The Oracle VM Server is removed from the server pool.

12. Repeat this procedure for each compute node that requires upgrade from Release 1.0.2 to Release 1.1.3.

Warning

The server pool master server must be the last Oracle VM Server to go through this upgrade procedure.

13. When all compute nodes have been upgraded and rebooted successfully, acknowledge all events for the Oracle VM Servers and the *Rack1-Repository*.

    Information on acknowledging events is provided in the Oracle VM User's Guide section entitled Acknowledging Events/Errors.

14. Remove the stale LUN mappings left behind from the previous storage configuration.

    a. Go to the **Storage** tab of the Oracle VM Manager UI.

    b. In the navigation pane, select the SAN Server named *OVCA_ZFSSA_Rack1*.

    c. In the management pane, select the LUN or LUNs in *WARNING* status and click **Delete Physical Disk** in the toolbar.

15. Re-add any node that was removed from the *Rack1_ServerPool*. Make sure the *Rack1-Repository* is presented to all servers in the pool.

If no errors have occurred during this procedure, the software update to Release 1.1.3 is now complete.

7.7 Enabling Compute Node IPMI in Oracle VM Manager

Oracle PCA controller software Release 2.0.1 and later automatically configures the Intelligent Platform Management Interface (IPMI) for all compute nodes in Oracle VM Manager. Compute nodes provisioned and discovered in an earlier release do not have IPMI enabled in Oracle VM Manager. If compute nodes are shut down and have no IPMI configured, they must be powered on by pressing the power button on the front panel of the server. To be able to control the power status of compute nodes remotely, you may add their IPMI configuration manually by following the instructions in this section.

Note

The compute node's IPMI is the built-in Oracle Integrated Lights Out Manager (ILOM). Its default user name is `root` ; the default password is `Welcome1` . Look up the ILOM IP in the Hardware View of the Oracle PCA Dashboard: roll the mouse over the compute node in question and note the IP address in the pop-up window.

Configuring IPMI for Compute Nodes

1. Log in to the Oracle VM Manager UI.

   For details, see Section 4.2, “Logging in to the Oracle VM Manager Web UI”.

2. Go to the **Servers and VMs** tab.

3. In the navigation pane, select the compute node you wish to control through the IPMI.

   In the toolbar, click `Edit Server`.

4. In the Edit Server dialog box, go to the `IPMI` tab.

5. In the IPMI tab, make the following changes:

   - Select the *Enable Server IPMI* and *Change IPMI Password* check boxes.
   - Enter the administrative *User Name* and *Password* in the respective fields.
   - Enter the *IP Address* for the IPMI of this compute node.

6. Click OK. The compute node configuration in Oracle VM Manager is updated.

7.8 Enabling SNMP Server Monitoring

For troubleshooting or hardware monitoring, it may be useful to enable SNMP on the servers in your Oracle PCA. While the tools for SNMP are available, the protocol is not enabled by default. This section

explains how to enable SNMP with the standard Oracle Linux and additional Oracle PCA Management Information Bases (MIBs).

Enabling SNMP on the Management Nodes

1. Using SSH and an account with superuser privileges, log into the management node.

   Note

   The data center IP address used in this procedure is an example.

   ```
   # ssh root@10.100.1.101
   root@10.100.1.101's password:
   [root@ovcamn05r1 ~]#
   ```

2. Locate the necessary `rpm` packages in the mounted directory `/nfs/shared_storage/mgmt_image/Packages`, which resides in the `MGMT_ROOT` file system on the ZFS storage appliance. The following packages are part of the Oracle PCA ISO image, which is either already installed on your system or downloaded and unpacked on the shared storage during the update process:

   - ovca-snmp-0.9-3.el6.x86_64.rpm
   - net-snmp-libs-5.5-49.0.1.el6.x86_64.rpm
   - net-snmp-5.5-49.0.1.el6.x86_64.rpm
   - lm_sensors-libs-3.1.1-17.el6.x86_64.rpm
   - net-snmp-utils-5.5-49.0.1.el6.x86_64.rpm

3. Install these packages by running the following command:

   ```
   # rpm -ivh ovca-snmp-0.9-3.el6.x86_64.rpm net-snmp-libs-5.5-49.0.1.el6.x86_64.rpm \
   net-snmp-5.5-49.0.1.el6.x86_64.rpm lm_sensors-libs-3.1.1-17.el6.x86_64.rpm \
   net-snmp-utils-5.5-49.0.1.el6.x86_64.rpm
   ```

4. Create an SNMP configuration file: `/etc/snmp/snmpd.conf`.

   This is a standard sample configuration:

   ```
   rocommunity public
   syslocation MyDataCenter
   dlmod ovca /usr/lib64/ovca-snmp/ovca.so
   ```

5. Enable the `snmpd` service.

   ```
   # service snmpd start
   ```

6. If desired, enable the `snmpd` service on boot.

   ```
   # chkconfig snmpd on
   ```

7. Open the SNMP ports on the firewall.

   ```
   # iptables -I INPUT -p udp -m udp --dport 161 -j ACCEPT
   # iptables -I INPUT -p udp -m udp --dport 162 -j ACCEPT
   # iptables-save > /etc/sysconfig/iptables
   ```

   SNMP is now ready for use on this management node. Besides the standard Oracle Linux MIBs, these are also available:

- ORACLE-OVCA-MIB::ovcaVersion
- ORACLE-OVCA-MIB::ovcaSerial
- ORACLE-OVCA-MIB::ovcaType
- ORACLE-OVCA-MIB::ovcaStatus
- ORACLE-OVCA-MIB::nodeTable

Usage examples:

```
# snmpwalk  -v 1 -c public -O e 130.35.70.186 ORACLE-OVCA-MIB::ovcaVersion
# snmpwalk  -v 1 -c public -O e 130.35.70.111 ORACLE-OVCA-MIB::ovcaStatus
# snmpwalk  -v 1 -c public -O e 130.35.70.111 ORACLE-OVCA-MIB::nodeTable
```

8. Repeat this procedure on the second management node.

Enabling SNMP on the Compute Nodes

Note

On Oracle PCA compute nodes, `net-snmp`, `net-snmp-utils` and `net-snmp-libs` are already installed at the factory, but the SNMP service is not enabled or configured.

1. Using SSH and an account with superuser privileges, log into the compute node. It can be accessed through the appliance internal management network.

```
ssh root@192.168.4.5
root@192.168.4.5's password:
[root@ovcacn27r1 ~]#
```

2. Create an SNMP configuration file: `/etc/snmp/snmpd.conf` and make sure this line is included:

```
rocommunity public
```

3. Enable the `snmpd` service.

```
# service snmpd start
```

SNMP is now ready for use on this compute node.

4. If desired, enable the `snmpd` service on boot.

```
# chkconfig snmpd on
```

5. Repeat this procedure on all other compute nodes installed in your Oracle PCA environment.

7.9 Changing Oracle WebLogic Server Passwords

The Oracle PCA Dashboard does not have the functionality to manage the Oracle WebLogic Server passwords. However, for security reasons, it is advisable that you replace the default password. It is not technically required to use the same system password you set in the Dashboard. To make sure that the appliance software can continue to connect to Oracle WebLogic Server, you must configure the new password in two locations:

- the Oracle WebLogic Server Administration Console

- the Oracle PCA Password Manager, or Wallet

Changing the Administrative Password in the Oracle WebLogic Server Administration Console

Note

This procedure uses the Administration Console, which is the web UI. If you prefer to use the Oracle WebLogic Scripting Tool (WLST) instead, follow the instructions in the WLST documentation, Chapter 6, “Configuring Existing WebLogic Domains”: Changing a Password.

1. In your browser, enter the address `https://IP:7002/console` .

   In the address, *IP* refers to the IP address that you have configured for the management nodes during installation.

2. Log in with the user name `weblogic` and the global password. [1]

3. Change the password of the user named "`weblogic`" by executing the steps described in the section “Change user passwords” in the Oracle WebLogic Server Administration Console Online Help.

Note

The Oracle WebLogic Server instances running on both management nodes work together as a cluster. The password update on one instance is automatically replicated on the other.

Caution

Keep Oracle WebLogic Server and the Password Manager in sync. Otherwise, Oracle PCA connections to Oracle WebLogic Server will fail.

Changing the Oracle WebLogic Server Password in the Oracle Private Cloud Appliance Password Manager (Wallet)

1. Using SSH and an account with superuser privileges, log into the master management node.

Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Set a new Wallet entry for the `weblogic` user.

```
# pca-update-password -t wls
```

Note

Password changes through the Dashboard, which apply to all other system configuration components of Oracle PCA, are described in Section 2.4, “Network Settings”.

7.10 Replacing Default Passwords Manually

[1] The factory-default global password is `Welcome1` .

The password update functionality in the Oracle PCA Dashboard is restricted to these system configuration components: the Dashboard UI itself, Oracle VM Manager, the root user account on both management nodes, and the ovs user account of the Oracle VM mySQL database. However, for security reasons, you may prefer to replace the default password for certain user accounts and components. It is not technically required to use the same password you set in the Oracle PCA Dashboard. To make sure that the appliance software can continue to connect to all its components, you must configure the new password in two locations:

- the component's access control mechanism; typically a UI or CLI
- the Oracle PCA Password Manager (Wallet)

Caution

To reset passwords that are out of sync for certain components, for example after restoring an appliance software and configuration backup, execute the steps below that apply to the Wallet entries.

Changing the Administrative Password for an Oracle Private Cloud Appliance Component

1. Log in to the system component for which you want to change the administrative password. [2]
2. Change the administrative password.

   For instructions, refer to the applicable product documentation. Links to related documentation are listed in the Preface.

   Caution

   Keep the Password Manager in sync with each of your password changes. Otherwise, Oracle PCA may fail to connect to its components.

3. Using SSH and an account with superuser privileges, log into the master management node.

   Note

   The data center IP address used in this procedure is an example.

   ```
   # ssh root@10.100.1.101
   root@10.100.1.101's password:
   [root@ovcamn05r1 ~]#
   ```

4. Set a new Wallet entry for the corresponding component or user account.

   ```
   # pca-update-password -t <password_type>
   ```

   Valid password types are: `nm2`, `xsigo`, `ilom`, `db`, `agent`, `pca`, `mgmt`, `wallet`, `zfs`, `pdu`, `wls`, `system`, `opus`, `ovmm`, `xms`.

   Note

   Password changes through the Dashboard, which apply to all other system configuration components of Oracle PCA, are described in Section 2.4, “Network Settings”.

[2] The factory-default global password is `Welcome1` .

7.11 Using a Custom CA Certificate for SSL Encryption

By default, Oracle PCA and Oracle VM Manager use a self-signed SSL certificate. While it serves to provide SSL encryption for all HTTP traffic, it is recommended that you obtain and install your own custom trusted certificate from a well-known and recognized Certificate Authority (CA).

Both the Oracle PCA Dashboard and the Oracle VM Manager web interface run on Oracle WebLogic Server. Oracle WebLogic Server provides the functionality to update the digital certificate and keystore. To add your own trusted CA certificate and keystore, see the procedures set out in the Oracle WebLogic documentation:

- Configuring Keystores: http://docs.oracle.com/cd/E17904_01/apirefs.1111/e13952/taskhelp/security/ConfigureKeystoresAndSSL.html
- Configuring Identity and Trust: http://docs.oracle.com/cd/E23943_01/web.1111/e13707/identity_trust.htm

Caution

Each Oracle PCA management node runs its own instance of Oracle WebLogic Server. You must apply the same changes separately to the master and standby management node.

Accessing the Oracle WebLogic Server Administration Console

1. In your browser, enter the address `https://IP:7002/console` .

   In the address, *IP* refers to the IP address that you have configured for each of the management nodes during installation.

2. Log in with the user name `weblogic` and the global password. [1]

   If you have already changed the default password, enter the password you configured by following the instructions in Section 7.9, “Changing Oracle WebLogic Server Passwords”.

7.12 A Compute Node Fails to Complete Provisioning

Compute node provisioning is a complex orchestrated process involving various configuration and installation steps and several reboots. Due to connectivity fluctuations, timing issues or other unexpected events, a compute node may become stuck in an intermittent state or go into error status. The solution is to reprovision the compute node.

Warning

Reprovisioning is to be applied *only* to compute nodes that fail to complete provisioning.

For correctly provisioned and running compute nodes, reprovisioning functionality is blocked in order to prevent incorrect use that could lock compute nodes out of the environment permanently or otherwise cause loss of functionality or data corruption.

Reprovisioning a Compute Node when Provisioning Fails

1. Log in to the Oracle PCA Dashboard.
2. Go to the **Hardware View** tab.
3. Roll over the compute nodes in Warning or Error status.

A pop-up window displays a summary of configuration and status information.

Figure 7.1 Compute Node Information and Reprovision Button in Hardware View

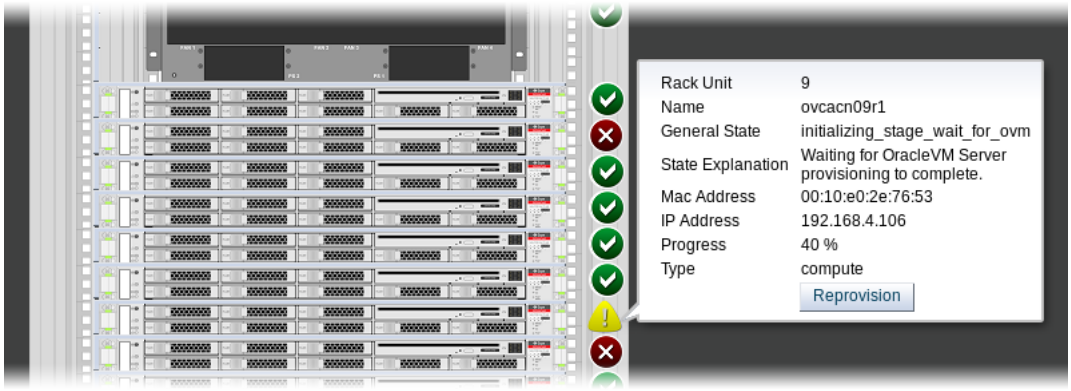

4. If the compute node provisioning is incomplete and the server is in error status or stuck in an intermittent state for several hours, click the **Reprovision** button in the pop-up window.

5. When the confirmation dialog box appears, click OK to start reprovisioning the compute node.

If compute node provisioning should fail after the server was added to the Oracle VM server pool, additional recovery steps could be required. The cleanup mechanism associated with reprovisioning may be unable to remove the compute node from the Oracle VM configuration. For example, when a server is in locked state or owns the server pool master role, it must be unconfigured manually. In this case you need to perform operations in Oracle VM Manager that are otherwise not permitted. You may also need to power on the compute node manually.

Removing a Compute Node from the Oracle VM Configuration

1. Log into the Oracle VM Manager user interface.

   For detailed instructions, see Section 4.2, “Logging in to the Oracle VM Manager Web UI”.

2. Go to the **Servers and VMs** tab and verify that the server pool named `Rack1_ServerPool` does indeed contain the compute node that fails to provision correctly.

3. If the compute node is locked due to a running job, abort it in the **Jobs** tab of Oracle VM Manager.

   Detailed information about the use of jobs in Oracle VM can be found in the Oracle VM User's Guide. Refer to the sections entitled Jobs Tab and Working with the Jobs Framework.

4. Remove the compute node from the Oracle VM server pool.

   Refer to the Oracle VM User's Guide and follow the instructions in the section entitled Removing an Oracle VM Server from a Server Pool.

5. Delete the compute node from Oracle VM Manager.

   Refer to the Oracle VM User's Guide and follow the instructions in the section entitled Deleting Oracle VM Servers from Oracle VM Manager.

When the failing compute node has been removed from the Oracle VM configuration, return to the Oracle PCA Dashboard, to reprovision it. If the compute node is powered off and reprovisioning cannot be started, power on the server manually.

7.13 Oracle VM Server Pool Is Offline After Network Services Restart

When network services are restarted on the master management node, the connection to the Oracle VM management network (`bond0`) is lost. By design, the `bond0` interface is not brought up automatically on boot, so that the virtual IP of the management cluster can be configured on the correct node, depending on which management node assumes the master role. While the master management node is disconnected from the Oracle VM management network, the Oracle VM Manager user interface reports that the compute nodes in the server pool are offline.

The management node that becomes the master, runs the Oracle VM services necessary to bring up the `bond0` interface and configure the virtual IP within a few minutes. It is expected that the compute nodes in the Oracle VM server pool return to their normal online status in the Oracle VM Manager user interface. If the master management node does not reconnect automatically to the Oracle VM management network, bring the `bond0` interface up manually from the Oracle Linux shell.

Warning

Execute this procedure *ONLY* when so instructed by Oracle Support. This should only be necessary in rare situations where the master management node fails to connect automatically. You should never manually disconnect or restart networking on any node.

Manually Reconnecting the Master Management Node to the Oracle VM Management Network

1. Using SSH and an account with superuser privileges, log into the disconnected master management node on the appliance management network.

```
# ssh root@192.168.4.3
root@192.168.4.3's password:
[root@ovcamn05r1 ~]#
```

2. Check the configuration of the `bond0` interface.

   If the interface is down, the console output looks similar to this:

```
# ifconfig bond0
bond0     Link encap:Ethernet  HWaddr 00:13:97:4E:B0:02
          BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 B)  TX bytes:0 (0.0 B)
```

3. Bring the `bond0` interface up.

```
# ifconfig bond0 up
```

4. Check the configuration of the `bond0` interface again.

   When the interface reconnects successfully to the Oracle VM management network, the console output looks similar to this:

```
# ifconfig bond0
bond0     Link encap:Ethernet  HWaddr 00:13:97:4E:B0:02
          inet addr:192.168.140.4  Bcast:192.168.140.255  Mask:255.255.255.0
          inet6 addr: fe80::213:97ff:fe4e:b002/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
          RX packets:62191 errors:0 dropped:0 overruns:0 frame:0
          TX packets:9183 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
```

```
          RX bytes:4539474 (4.33 MB)  TX bytes:1853641 (1.77 MB)
```

Index

A

appliance
- architecture, 1
- backup, 17
- hardware, 1, 2
- high availability, 16
- overview, 1, 1
- provisioning, 1, 14
- software, 1, 8
- software update, 18

ASR, 69
- activate assets, 75
- install ASR Manager, 71
- introduction, 69
- prerequisites, 70

ASR Manager
- configuring failover, 73
- starting the service, 74

authentication, 27

Auto Service Request, 69
- activate assets, 75
- configuring ASR Manager failover, 73
- install ASR Manager, 71
- introduction, 69
- prerequisites, 70
- starting ASR Manager, 74

B

backup, 17, 21

C

compute node
- enable IPMI, 118
- enable support tools for release 1.0.2, 113
- failover, 16
- provisioning, 14
- server information, 59
- upgrade 1.0.2 to 1.1.1 or later, 115
- upgrade Oracle VM Server, 33
- upgrade virtualization platform, 33

compute nodes
- hardware, 4

configuration, 21, 25, 25, 25, 25, 25
- logging, 112
- proxy, 111

CRU servicing
- NM2-36P Sun Datacenter InfiniBand Expansion Switch, 106
- Oracle Fabric Interconnect F1-15, 107
- Oracle Server X5-2, 92
- Oracle Switch ES1-24, 105
- Oracle ZFS Storage Appliance ZS3-ES, 99
- rack parts, 90
- Sun Server X3-2, 97
- Sun Server X4-2, 94
- Sun ZFS Storage Appliance 7320, 102

D

Dashboard, 21, 25, 27
- Hardware View, 22
- login, 22
- Network View, 25
- software, 8

database
- failover, 16

databases
- software, 9

DNS, 25

E

electrostatic discharge safety, 89

ES1-24 switch
- component servicing instructions, 105

F

Fabric Interconnect F1-15
- component servicing instructions, 107
- replaceable components, 89

failover
- compute node, 16
- database, 16
- management node, 16
- networking, 17
- storage, 17

H

hardware, 2
- compute nodes, 4
- identifying, 22
- management nodes, 3
- monitoring, 22
- networking, 6
- status, 22
- storage, 5
- view, 22

high availability, 16

I

InfiniBand switch
- component servicing instructions, 106
- replaceable components, 88

initial setup, 21, 25, 28

IPMI

enable, 118

J

jobs and events, 67

L

logging
configure, 112

M

management node, 25
failover, 16
provisioning, 14
management nodes
hardware, 3
monitoring
enabling SNMP, 118
hardware, 21, 22, 59
network, 25
Oracle VM, 59
virtual machine, 59

N

network
configuration, 21
hardware, 25
identifying, 25
monitoring, 21, 25
view, 25
Network Settings, 25
networking
failover, 17
hardware, 6
proxy, 111
using VLANs, 112
NM2-36P InfiniBand switch
component servicing instructions, 106
replaceable components, 88
NTP, 25

O

operating systems
software, 9
Oracle Server X5-2, 4
component servicing instructions, 92
replaceable components, 83
Oracle VM, 57
Events, 67
health, 59
Jobs, 67
limitations, 58
login, 58
monitoring, 59
networking, 64
repositories, 62
Servers and VMs, 59
storage, 66
tagging, 67
Tools and Resources, 67
Oracle VM Manager
adding expansion nodes, 15
server pool, 15
software, 9
verify tuning settings, 114
Oracle ZFS Storage Appliance ZS3-ES, 5
component servicing instructions, 99

P

password, 27
password manager, 8
change password, 122
change WebLogic password, 120
power
enabling IPMI, 118
power off procedure
servicing, 90
provisioning, 14
compute node discovery, 14
expansion node, 15
initialization, 14
server pool configuration, 15
proxy
configure, 111
public network, 25

R

rack
component servicing instructions, 90
replaceable components, 82
replaceable components
NM2-36P Sun Datacenter InfiniBand Expansion Switch, 88
Oracle Fabric Interconnect F1-15, 89
Oracle Server X5-2, 83
Oracle Switch ES1-24, 88
Oracle ZFS Storage Appliance ZS3-ES, 85
rack infrastructure, 82
Sun Server X3-2, 84
Sun Server X4-2, 84
Sun ZFS Storage Appliance 7320, 86
reset, 25, 27

S

safety
electrostatic discharge, 89
service precautions, 89

service, 81
  Auto Service Request, 69
  electrostatic discharge, 89
  preparations, 89
  replaceable components, 81
  safety precautions, 89
  servicing Fabric Interconnect parts, 107
  servicing NM2-36P Sun Datacenter InfiniBand Expansion Switch parts, 106
  servicing Oracle Server X5-2 parts, 92
  servicing Oracle Switch ES1-24 parts, 105
  servicing Oracle ZFS Storage Appliance ZS3-ES parts, 99
  servicing rack parts, 90
  servicing Sun Server X3-2 parts, 97
  servicing Sun Server X4-2 parts, 94
  servicing Sun ZFS Storage Appliance 7320 parts, 102
servicing the system
  powering down, 90
SNMP
  enable, 118
software, 8
  Dashboard, 8
  databases, 9
  operating systems, 9
  Oracle VM Manager, 9
  wallet, 8
software update, 18, 21, 28
  upgrade existing compute node to release 1.1.1 or later, 115
  upgrade existing compute node with new Oracle VM Server, 33
  verify Oracle VM Manager tuning settings, 114
storage
  adding, 66
  configuration, 66
  failover, 17
  hardware, 5
Storage Appliance
  component servicing instructions, 99, 102
  replaceable components, 85, 86
Sun Server X3-2, 4
  component servicing instructions, 97
  replaceable components, 84
Sun Server X4-2, 4
  component servicing instructions, 94
  replaceable components, 84
Sun ZFS Storage Appliance 7320
  component servicing instructions, 102
support tools
  enabling on release 1.0.2, 113
switch ES1-24
  component servicing instructions, 105
  replaceable components, 88

T

tagging resources, 67

U

update, 21, 28
  upgrade compute node from 1.0.2 to 1.1.1 or later, 115
  upgrade Oracle VM Server, 33
  upgrade virtualization platform, 33
  verify Oracle VM Manager tuning settings, 114

V

virtual machine
  assemblies, 62
  clone, 59
  clone customizer, 62
  console, 59
  create, 59
  delete, 59
  import, 59
  installation media, 62
  ISOs, 62
  management, 57, 59
  messaging, 59
  migrate, 59
  resources, 62
  templates, 62
  virtual disks, 62
  VNICs, 64
VLAN
  enabling VLAN traffic, 112

W

wallet
  change password, 121
  software, 8
  WebLogic password, 120
WebLogic
  change password, 120

Z

ZFS Storage Appliance
  component servicing instructions, 99, 102
  replaceable components, 85, 86